

The Capacity of Finite-State Channels in the High-Noise Regime

Henry D. Pfister*

Abstract

This paper considers the derivative of the entropy rate of a hidden Markov process with respect to the observation probabilities. The main result is a compact formula for the derivative that can be evaluated easily using Monte Carlo methods. It is applied to the problem of computing the capacity of a finite-state channel (FSC) and, in the high-noise regime, the formula has a simple closed-form expression that enables series expansion of the capacity of a FSC. This expansion is evaluated for a binary-symmetric channel under a (0,1) run-length limited constraint and an intersymbol-interference channel with Gaussian noise.

1 Introduction

1.1 The Hidden Markov Process

A hidden Markov process (HMP) is a discrete-time finite-state Markov chain (FSMC) observed through a memoryless channel. The HMP has become ubiquitous in statistics, computer science, and electrical engineering because it approximates many processes well using a dependency structure that leads to many efficient algorithms. While the roots of the HMP lie in the “grouped Markov chains” of Harris [21] and the “functions of a finite-state Markov chain” of Blackwell [8], the HMP first appears (in full generality) as the output process of a finite-state channel (FSC) [9]. The statistical inference algorithm of Baum and Petrie [5], however, cemented the HMP’s place in history and is responsible for great advances in fields such as speech recognition and biological sequence analysis [23, 25]. An exceptional survey of HMPs, by Ephraim and Merhav, gives a nice summary of what is known in this area [13].

Definition 1.1. Let \mathcal{Q} be the state set of an irreducible aperiodic FSMC $\{Q_t\}_{t \in \mathbb{Z}}$ with state transition matrix P and define

$$p_{ij} \triangleq [P]_{i,j} = \Pr(Q_{t+1} = j | Q_t = i)$$

for $i, j \in \mathcal{Q}$. Let \mathcal{Y} be a finite set of possible observations and $\{Y_t\}_{t \in \mathbb{Z}}$ be the stochastic process where $Y_t \in \mathcal{Y}$ is generated by the transition from Q_t to Q_{t+1} . The distribution of the observation conditioned on the FSMC transition¹ is given by

$$h_{ij}(y) \triangleq \begin{cases} \Pr(Y_t = y | Q_t = i, Q_{t+1} = j) & \text{if } (i, j) \in \mathcal{V} \\ 0 & \text{otherwise} \end{cases}$$

for $i, j \in \mathcal{Q}$, where $\mathcal{V} = \{(i, j) \in \mathcal{Q} \times \mathcal{Q} | p_{ij} > 0\}$ is the set of valid transitions. The ergodic process $\{Y_t\}_{t \in \mathbb{Z}}$ is called a *hidden Markov process*. With proper initialization, the process is also stationary.

Although the notation of this paper assumes that \mathcal{Y} is a finite set, many results remain correct when $\mathcal{Y} = \mathbb{R}$ if $h_{ij}(y)$ is assumed to be a continuous p.d.f. and sums over \mathcal{Y} are converted to integrals over \mathbb{R} .

*Henry Pfister is with the Electrical and Computer Engineering Department of Texas A&M University (e-mail: hpfister@tamu.edu). His research was supported in part by the National Science Foundation under Grant No. 074740.

¹In general, HMPs are defined by noisy observations of the FSMC states (rather than the transitions). This paper uses the “transition observation” model instead because of its natural connection with finite-state channels. Moreover, any random process that can be represented by the “transition observation” HMP model with M states can also be represented by the “state observation” model with M^2 states.

1.2 The Entropy Rate

The entropy rate of a stationary stochastic process $\{Y_t\}_{t \in \mathbb{Z}}$ is defined to be

$$H(\mathcal{Y}) \triangleq \lim_{n \rightarrow \infty} \frac{1}{n} H(Y_1, \dots, Y_n),$$

where $H(Y_1) \triangleq -E[\ln \Pr(Y_1)]$ is the entropy of the random variable (r.v.) Y_1 and the limit exists and is finite if $H(Y_1) < \infty$ [11]. Computing the exact entropy rate of an HMP in closed form appears to be difficult, however. In [8], Blackwell states

“In this paper we study the entropy of the $\{y_n\}$ [hidden Markov] process; our result suggests that this entropy is intrinsically a complicated function of [the parameters of the hidden Markov process] M and Φ .”

On the other hand, the Shannon-McMillan-Breiman Theorem shows that the empirical entropy rate $-\frac{1}{n} \ln \Pr(y_1^n)$ converges almost surely to the entropy rate $H(\mathcal{Y})$ (in nats) as $n \rightarrow \infty$. Therefore, simulation based (i.e., Monte Carlo) approaches work well in many cases [30, 17, 1, 38, 37, 3, 2].

Other early work related to the entropy rate of HMPs can be found in [7, 36, 39, 35]. Recently, interest in HMPs has surged and there have been a large number of papers discussing the entropy rate of HMPs. These range from bounds [37, 31, 32] to establishing the analyticity of the entropy rate [18] to computing series expansions of the entropy rate [44, 12, 20].

1.3 The Finite-State Channel

The work in this paper is largely motivated by the analysis of a class of time-varying channels known as FSCs. An FSC is a discrete-time channel where the distribution of the channel output depends on both the channel input and the underlying channel state [16]. This allows the channel output to depend implicitly on previous inputs and outputs via the channel state. In practice, there are three types of channel variation which FSCs are typically used to model. A *flat fading* channel is a time-varying channel whose state is independent of the channel inputs. An *intersymbol-interference* (ISI) channel is a time-varying channel whose state is a deterministic function of the previous channel inputs. Channels which exhibit both fading and ISI can also be modeled, and their state is a stochastic function of the previous channel inputs. An *indecomposable FSC* is, roughly speaking, a FSC where the effect of the initial state decays with time. The output process of an indecomposable FSC with an ergodic Markov input is an HMP.

Consider an indecomposable FSC with state set \mathcal{S} , finite input alphabet \mathcal{X} , and output alphabet \mathcal{Y} . The channel is defined by its input-output state-transition probability $W(y, s' | x, s)$, which is defined for all $x \in \mathcal{X}$, $y \in \mathcal{Y}$, and $s, s' \in \mathcal{S}$. Using this notation, $W(y, s' | x, s)$ is the conditional probability that the channel output is y and the new channel state is s' given that the channel input was x and the initial state was s . The n -step transition probability for a sequence of n channel uses (with input x_1^n and output y_1^n) is given by

$$\Pr(Y_1^n = y_1^n | X_1^n = x_1^n) = \sum_{s_1^{n+1} \in \mathcal{S}^{n+1}} \Pr(S_1 = s_1) \prod_{t=1}^n W(y_t, s_{t+1} | x_t, s_t).$$

When $\mathcal{Y} = \mathbb{R}$, we will also use $W(y, s' | x, s)$ to represent a conditional probability density function for the channel outputs.

The achievable information rate of an FSC with Markov inputs is intimately related to the entropy rate of an HMP [1, 38, 24, 2, 42, 22]. Computing this entropy rate exactly is usually quite difficult, and often the main obstacle in the computation of achievable rates.

1.4 Main Results

The main result of this paper, given in Theorem 3.2, is a compact formula for the derivative, with respect to the observation probability $h_{ij}(y)$, of the entropy rate of a general HMP. A Monte Carlo

estimator for this derivative follows easily because the formula is an expectation over distributions that are relatively easy to sample. The formula is also amenable to analysis in some asymptotic regimes. In particular, Theorem 3.6 derives a simple formula for the first two non-trivial terms in the expansion of the entropy rate in the high-noise regime.

In Section 4, this derivative formula also allows one to consider the derivative of achievable information rates for FSCs. For example, a closed-form expression for the capacity of a BSC under a (0,1) RLL constraint is derived in the high-noise limit. Section 2 provides the mathematical background necessary for the later sections.

2 Mathematical Background

2.1 Notation

Calligraphic letters are used to denote sets (e.g., $\mathcal{Q}, \mathcal{Y}, \mathcal{V}$) and $\mathbb{1}_{\mathcal{Y}}(\cdot)$ is the indicator function of the set \mathcal{Y} . Capital letters are used to denote random variables (e.g., Q_t, Y_t) and matrices (e.g., M, P). Lower-case letters are used to represent realizations of random variables (e.g., q_t, y_t), column vectors (e.g., π, α, β, u, v), and indices (e.g., i, j, k, l). The i -th element of the vector π is denoted $\pi(i)$.

The following sets will also be used: $\mathbb{R}_+ = \{a \in \mathbb{R} \mid a > 0\}$, $\mathcal{A} = \mathbb{R}^{|\mathcal{Q}|}$, $\mathcal{A}_\delta = \{u \in \mathcal{A} \mid u(q) > \delta, q \in \mathcal{Q}\}$, $\mathcal{P} = \{u \in \mathcal{A} \mid \sum_q u(q) = 1\}$, and $\mathcal{P}_\delta = \mathcal{A}_\delta \cap \mathcal{P}$. We note that the symbols $\pi, \alpha_t \in \mathcal{P}_0$ are used interchangeably to denote distributions over \mathcal{Q} and $|\mathcal{Q}|$ -dimensional column vectors (e.g., $\pi^T P = \pi^T$). The standard p -norm of the vector u is denoted by $\|u\|_p \triangleq (\sum_i |u(i)|^p)^{1/p}$ and the induced matrix norm is $\|M\|_p \triangleq \sup_{\|u\|_p=1} \|Mu\|_p$.

2.2 The Forward-Backward Algorithm

One of the primary reasons for the popularity of HMPs is that the forward and backward state estimation problems have a simple recursive structure. Let us assume that the Markov chain $\{Q_t\}_{t \in \mathbb{Z}}$ is stationary and that $\pi \in \mathcal{P}_0$ is the unique stationary distribution that satisfies $\pi^T P = \pi^T$. For a length- n block, let the forward state probability $\alpha_t \in \mathcal{P}$ and the backward state probability $\beta_t \in \mathcal{A}$ be defined by

$$\begin{aligned}\alpha_t(i) &\triangleq \Pr(Q_t = i \mid Y_1^{t-1} = y_1^{t-1}) \\ \beta_t(j) &\triangleq \frac{1}{\pi(j)} \Pr(Q_t = j \mid Y_t^n = y_t^n)\end{aligned}$$

for $i, j \in \mathcal{Q}$. These definitions lead naturally to the recursions

$$\begin{aligned}\alpha_{t+1}(j) &= \frac{1}{\psi_{t+1}} \sum_{i \in \mathcal{Q}} \alpha_t(i) p_{ij} h_{ij}(y_t) \\ \beta_{t-1}(i) &= \frac{1}{\phi_{t-1}} \sum_{j \in \mathcal{Q}} \beta_t(j) p_{ij} h_{ij}(y_{t-1})\end{aligned}$$

for $i, j \in \mathcal{Q}$, where ψ_{t+1} is chosen so that $\sum_{i \in \mathcal{Q}} \alpha_{t+1}(i) = 1$ and ϕ_{t-1} is chosen² so that $\sum_{j \in \mathcal{Q}} \pi(j) \beta_{t-1}(j) = 1$. It is worth noting that $\psi_{t+1} = \Pr(Y_t = y_t \mid Y_1^{t-1} = y_1^{t-1})$ and therefore we find that

$$-\frac{1}{n} \sum_{t=1}^n \ln \psi_{t+1} = -\frac{1}{n} \ln \Pr(Y_1^n = y_1^n) \xrightarrow{a.s.} H(\mathcal{Y}) \text{ nats.}$$

This simple connection between the forward recursion and the entropy rate implies a simple Monte Carlo approach to estimating the achievable information rates of FSCs [1, 38, 37, 3, 2].

²We believe this normalization for $\beta_{i-1}(q)$ is new and it appears to be the natural choice for the problem considered in this paper (and perhaps in general).

2.3 The Matrix Perspective

2.3.1 The Forward-Backward Algorithm

In this section, we review a natural connection between the product of random matrices and the forward-backward recursions. This connection is interesting in its own right, but will also be very helpful in understanding the results of later sections.

Definition 2.1. For any $y \in \mathcal{Y}$, the *transition-observation probability matrix*, $M(y)$, is a $|\mathcal{Q}| \times |\mathcal{Q}|$ matrix defined by

$$[M(y)]_{ij} \triangleq \Pr(Y_t = y, Q_{t+1} = j \mid Q_t = i) = p_{ij} h_{ij}(y). \quad (2.1)$$

These matrices behave similarly to transition probability matrices because their sequential products compute the n -step transition observation probabilities of the form,

$$[M(y_t)M(y_{t+1}) \dots M(y_{t+k})]_{ij} = \Pr(Y_t^{t+k} = y_t^{t+k}, Q_{t+k+1} = j \mid Q_t = i).$$

This means that we can write $\Pr(Y_1^n = y_1^n)$ as the matrix product³

$$\Pr(Y_1^n = y_1^n) = \pi^T \left(\prod_{t=1}^n M(y_t) \right) \mathbf{1}, \quad (2.2)$$

where $\mathbf{1}$ is a $|\mathcal{Q}|$ -dimensional column vector of ones. When $\mathcal{Y} = \mathbb{R}$, the above expressions are understood to be probability density functions with respect to the observations and the joint probability becomes the joint density.

Likewise, the forward/backward recursions can be written in matrix form as

$$\alpha_{t+1}^T = \frac{\alpha_t^T M(y_t)}{\alpha_t^T M(y_t) \mathbf{1}} \quad \beta_{t-1} = \frac{M(y_{t-1}) \beta_t}{\pi^T M(y_{t-1}) \beta_t}$$

where $\pi^T \mathbf{1} = 1$, $\alpha_{t+1}^T \mathbf{1} = 1$, and $\pi^T \beta_{t-1} = 1$. We will also make use of the shorthand notation

$$M(y_k^l) \triangleq \prod_{t=k}^l M(y_t).$$

2.3.2 Contraction Coefficients

This section summarizes some standard results on the contractive properties of positive matrices and their connections to HMPs. More details can be found in [40, 27, 26].

Definition 2.2. For any two vectors $u, v \in \mathcal{A}_0$, the *Hilbert projective metric* is

$$d(u, v) \triangleq \ln \frac{\max_i (u(i)/v(i))}{\min_j (u(j)/v(j))} = \ln \max_{i,j} \frac{u(i)v(j)}{v(i)u(j)} = -\ln \min_{i,j} \frac{u(i)v(j)}{v(i)u(j)}.$$

It is metric on $\mathcal{A}_0 \setminus \sim$ where \sim is the equivalence relation with $u \sim v$ if $au = v$ for some $a \in \mathbb{R}_+$.

Proposition 2.3. For $u, v, w \in \mathcal{A}_0$ such that $w^T u = w^T v$, the Hilbert projective metric characterizes the element-wise relative distance between two vectors in the sense that, for any $i \in \mathcal{Q}$,

$$\begin{aligned} d_M(u(i), v(i)) &\triangleq \frac{|u(i) - v(i)|}{\max(u(i), v(i))} \leq 1 - e^{-d(u, v)} \leq d(u, v) \\ d_m(u(i), v(i)) &\triangleq \frac{|u(i) - v(i)|}{\min(u(i), v(i))} \leq e^{d(u, v)} - 1 \stackrel{“d(u, v) \leq 1”}{\leq} 2d(u, v), \end{aligned}$$

where d_M is a metric on \mathbb{R}_+ and d_m is a semi-metric on \mathbb{R}_+ (i.e., the triangle inequality does not hold).

³Since matrix multiplication is not commutative, we use the convention that $\prod_{t=1}^n M(y_t) = M(y_1)M(y_2) \dots M(y_n)$.

Proof. If $u(k) \geq v(k)$, then we have

$$u(k)e^{-d(u,v)} = u(k) \min_j \frac{v(j)}{u(j)} \min_i \frac{u(i)}{v(i)} \leq v(k) \min_i \frac{u(i)}{v(i)} \leq v(k),$$

where $\min_i \frac{u(i)}{v(i)} \leq 1$ because $w^T u = w^T v$. The stated results follow from $u(k) - v(k) \leq e^{d(u,v)} v(k) - v(k)$, $u(k) - v(k) \leq u(k) - u(k)e^{-d(u,v)}$, and simple bounds on e^x . Both distances are clearly symmetric and positive definite. The triangle inequality and other properties of d_M are discussed in [43]. \square

Lemma 2.4. *For any vectors $u, v, w \in \mathcal{A}_0$ such that $w^T u = w^T v$, we have*

$$\begin{aligned} \|u - v\|_1 &\leq (1 - e^{-d(u,v)}) \sum_{i \in \mathcal{Q}} \max(u(i), v(i)) \leq (\|u\|_1 + \|v\|_1) d(u, v) \\ \|u - v\|_1 &\leq (e^{d(u,v)} - 1) \sum_{i \in \mathcal{Q}} \min(u(i), v(i)) \leq (e^{d(u,v)} - 1) \min(\|u\|_1, \|v\|_1). \end{aligned}$$

Proof. The expressions follow from direct calculation of $\|u - v\|_1$ using the bounds in Proposition 2.3. \square

The following theorem of Birkhoff plays an important role in the remainder of this paper.

Theorem 2.5 ([40, Ch. 3]). *Consider any non-negative matrix M with at least one positive entry in every row and column. Then, for all $u, v \in \mathcal{A}_0$, we have*

$$d(Mu, Mv) \leq \tau(M) d(u, v)$$

where $\tau(M) \triangleq \frac{1 - \phi(M)^{1/2}}{1 + \phi(M)^{1/2}} = \tau(M^T) \leq 1$ is the Birkhoff contraction coefficient and

$$\phi(M) = \min_{i,j,k,l} \frac{[M]_{ik} [M]_{jl}}{[M]_{jk} [M]_{il}} \geq \left(\frac{\min_{i,j} [M]_{ij}}{\max_{i,j} [M]_{ij}} \right)^2. \quad (2.3)$$

The following results connect our HMP definition with Birkhoff's contraction coefficients. An FSMC that is irreducible and aperiodic is called *primitive*. Since the underlying Markov chain is primitive, the matrix P must have at least one non-zero entry in each row and column.

Condition 2.6. For some $\delta \geq 0$, the joint probability of every valid transition and output is greater than δ . In other words, this means that $p_{ij} h_{ij}(y) > \delta \geq 0$ for all $(i, j) \in \mathcal{V}$ and $y \in \mathcal{Y}$.

Under Condition 2.6, the matrix $M(y)$ has exactly the same pattern of zero/non-zero entries as P for all $y \in \mathcal{Y}$. Since P is transition matrix for an ergodic Markov chain, one finds that $M(y)$ must also have at least one non-zero entry in each row and column for all $y \in \mathcal{Y}$. Therefore, $\tau(M(y)) \leq 1$ for all $y \in \mathcal{Y}$.

Definition 2.7. An HMP is said to be (ϵ, k) -*primitive* if $\min_{i,j} [M(y_1^k)]_{ij} > k\epsilon$ for all $y_1^k \in \mathcal{Y}$. This gives a uniform lower bound on the probability that a k -step transition of the HMP simultaneously moves between any two states and generates any output sequence y_1^k . An HMP is said to be ϵ -*primitive* if there exists a $k < \infty$ such it is (ϵ, k) -*primitive*.

Lemma 2.8. *An HMP is (ϵ, k) -primitive if it satisfies Condition 2.6 with $\delta \geq k^{1/k} \epsilon^{1/k}$ and P^k is a positive matrix. Moreover, this implies that $\pi(i) \geq k\epsilon$ (i.e., strictly positive) for all $i \in \mathcal{Q}$.*

Proof. First, we note that P^k positive implies there is a length- k path between any two states. Next, we write

$$\begin{aligned} [M(y_1^k)]_{q_1, q_{k+1}} &= \sum_{q_2, \dots, q_k \in \mathcal{Q}^{k-1}} \prod_{t=1}^k p_{q_t, q_{t+1}} h_{q_t, q_{t+1}}(y_t) \\ &> \sum_{q_2, \dots, q_k \in \mathcal{Q}^{k-1}} \prod_{t=1}^k \mathbb{1}_{\mathcal{V}}((q_t, q_{t+1})) \delta \\ &\stackrel{(a)}{\geq} \delta^k, \end{aligned}$$

where the last step follows from the fact that there is a length- k path between any two states. Since $\delta^k > k\epsilon$, we see that the HMP is (ϵ, k) -primitive according to Definition 2.7. Note that, for any $u \in \mathcal{A}_0$, we have

$$\sum_{i \in \mathcal{Q}} u(i) [M(y_1^k)]_{ij} \geq \left(\sum_{i \in \mathcal{Q}} u(i) \right) k\epsilon \geq \|u\|_1 k\epsilon \quad (2.4)$$

for $u \in \mathcal{A}_0$ implies that $\pi(i) \geq k\epsilon$ for all $i \in \mathcal{Q}$. \square

Lemma 2.9. *For any ϵ -primitive HMP, there exists a $k_0 < \infty$ such that, for all $y_1^k \in \mathcal{Y}^k$ and all $k \geq k_0$,*

$$\tau(M(y_1^k)) \leq e^{-2k_0 \lfloor k/k_0 \rfloor \epsilon}.$$

Proof. From Definition 2.7, we can assume that the HMP is (ϵ, k_0) -primitive. Using the bound (2.3), we see that

$$\phi(M(y_1^{k_0})) \geq \left(\frac{\min_{i,j} [M]_{ij}}{\max_{i,j} [M]_{ij}} \right)^2 \geq \left(\frac{k_0 \epsilon}{1} \right)^2$$

and

$$\tau(M(y_1^{k_0})) \leq \frac{1 - k_0 \epsilon}{1 + k_0 \epsilon} \leq e^{-2k_0 \epsilon}.$$

Since we can break any length k sequence into at least $\lfloor k/k_0 \rfloor$ length- k_0 pieces and $\tau(M(y)) \leq 1$ for the remaining pieces, we have $\tau(M(y_1^k)) \leq (e^{-2k_0 \epsilon})^{\lfloor k/k_0 \rfloor}$. \square

2.3.3 Lyapunov Exponents

Consider any stationary stochastic process, $\{Y_i\}_{i \in \mathbb{Z}}$, equipped with a function, $M(y)$, that maps each $y \in \mathcal{Y}$ to a matrix. Now, consider the limit

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \left\| u^T \prod_{i=1}^n M(Y_i) \right\|,$$

where u is any non-zero vector and $\|\cdot\|$ is any vector norm. Oseledec's multiplicative ergodic theorem says that this limit is deterministic for almost all realizations [34]. An earlier ergodic theorem of Furstenberg and Kesten [14] gives a nice proof that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \left\| \prod_{i=1}^n M(Y_i) \right\| \stackrel{a.s.}{=} \gamma_1,$$

where $\|\cdot\|$ is any matrix norm and γ_1 is known as the top Lyapunov exponent. The connection with entropy rate is given by the fact that, for an HMP, choosing $M(y)$ according to (2.1) implies that $H(\mathcal{Y}) = -\gamma_1$ [37, 22].

2.4 Stationary Measures

The forward and backward state probability vectors play a very important role in the analysis of HMPs. These vectors, $\alpha_i, \beta_i \in \mathcal{A}_0$, are themselves random variables which often have well-defined stationary distributions. To illustrate the mixing properties, we exploit the stationarity of the HMP and focus on time zero by defining the random variables

$$U_n(i) \triangleq \Pr(Q_0 = i | Y_{-n}^{-1})$$

$$V_n(i) \triangleq \frac{1}{\pi(i)} \Pr(Q_0 = i | Y_0^{n-1}).$$

It is worth noting that $U_n(i)$ is a deterministic function of y_{-n}^{-1} and $V_n(i)$ is a deterministic function of y_0^{n-1} . The following sufficient condition characterizes some of the HMPs that have stationary distributions.

Definition 2.10. An HMP is called *almost-surely mixing* if there exists a $C < \infty$, $\gamma < 1$, and $k < \infty$ such that

$$\begin{aligned}\Pr(d(U_m, U_n) > C\gamma^n) &\leq C\gamma^n \\ \Pr(d(V_m, V_n) > C\gamma^n) &\leq C\gamma^n\end{aligned}$$

for all $m \geq n + k \geq k + 1$. This implies that the forward and backward recursions both forget their initial conditions at an exponential rate that is uniform over all but an exponentially small set of received sequences.

Definition 2.11. An HMP is called *sample-path mixing* if there exists a $C < \infty$, $\gamma < 1$, and $k < \infty$ such that

$$\begin{aligned}d(U_m, U_n) &\leq C\gamma^n \\ d(V_m, V_n) &\leq C\gamma^n,\end{aligned}$$

for all $m \geq n + k \geq k + 1$ and all received sequences $y_{-m}^{m-1} \in \mathcal{Y}^{2m}$. This implies that the forward and backward recursions both forget their initial conditions at an exponential rate that is uniform over all received sequences. It is easy to see that sample-mixing implies almost-surely mixing.

Lemma 2.12. An (ϵ, k) -primitive HMP is sample-path mixing with $\gamma = e^{-2\epsilon}$ and $C = -2\ln(k\epsilon)\gamma^{-k}$.

Proof. For each y_{-n}^{-1} , the realization of $U_n(i)$ is given by

$$u_n(i) = \Pr(Q_0 = i | Y_{-n}^{-1} = y_{-n}^{-1}) = \left[\frac{\pi^T M(y_{-n}^{-1})}{\pi^T M(y_{-n}^{-1}) \mathbf{1}} \right]_i.$$

First, we let $w^T = \pi^T M(y_{-m}^{-n-1})$ and note that (2.4) implies that

$$d(w^T, \pi^T) = \ln \max_{i,j} \frac{w(i)\pi(j)}{\pi(i)w(j)} \leq \ln \max_{i,j} \frac{1}{\pi(i)w(j)} \leq \ln \left(\left(\frac{1}{k\epsilon} \right)^2 \right)$$

when $m \geq n + k$. Next, we use Theorem 2.5 and Lemma 2.9 to see that

$$\begin{aligned}d(u_m, u_n) &= d(w^T M(y_{-n}^{-1}), \pi^T M(y_{-n}^{-1})) d(w^T, \pi^T) \\ &\leq \tau(M(y_{-n}^{-1})) \ln((k\epsilon)^{-2}) \\ &\leq -2\ln(k\epsilon) e^{-2\lfloor n/k \rfloor k\epsilon}.\end{aligned}$$

This gives an exponential rate of $\gamma = e^{-2\epsilon}$ and $C = -2\ln(k\epsilon)\gamma^{-k}$ is chosen to handle the floor function and constant. For the backward recursion, the proof is identical except that the constant C is smaller by a factor of 2 because

$$d(M(y_n^{m-1})\mathbf{1}, \mathbf{1}) = \ln \max_{i,j} \frac{w(i)}{w(j)} \leq \ln \max_{i,j} \frac{\mathbf{1}^T \mathbf{1}}{\mathbf{1}^T \mathbf{1} k\epsilon} \leq \ln \left(\frac{1}{k\epsilon} \right).$$

□

Lemma 2.13. A $(0, k)$ -primitive HMP is almost-surely mixing for some $\gamma < 1$ and $C < \infty$ if

$$\max_{q \in \mathcal{Q}} E \left[\frac{\max_{i,j} [M(Y_1^k)]_{ij}}{\min_{i,j} [M(Y_1^k)]_{ij}} \middle| Q = q \right] < \infty.$$

In particular, this can be applied to HMPs with continuous observations.

Proof. This lemma follows, with slight modifications, from the arguments in [27]. Its proof is out of the scope of this work. □

Proposition 2.14. *The joint process $\{Q_t, \alpha_t\}_{t \in \mathbb{Z}}$ forms a Markov chain. If the HMP is almost-surely mixing, then the marginal distribution converges weakly to a unique stationary measure $\mu_q(A)$.*

Proof. One can see this is a Markov chain by considering the following method of generating the sequence. At each step, we first choose q_{t+1} according to $p_{q_t, q_{t+1}}$, then choose y_t according to $h_{q_t, q_{t+1}}(y_t)$, and finally compute $\alpha_{t+1}(\cdot)$ from $\alpha_t(\cdot)$ and y_t . In most cases, this Markov chain will not have a finite state-space because $\alpha_t(\cdot)$ may take uncountably many values. Of course, this process depends on the initialization of the first α_t but this dependence decays with time if the HMP is almost-surely mixing. For simplicity, one may assume the initialization $\alpha_1 = \pi$ is used.

To show that $\mu_q^{(t)}(A) \triangleq \Pr(Q_0 = q, U_t \in A)$ converges weakly to the probability measure $\mu_q(A)$ for all Borel subsets $A \subseteq \mathcal{P}_0$, we observe that $\mu_q^{(t)}(A)$ is a Cauchy sequence with respect to the Prohorov metric. This is sufficient because the Prohorov metric metrizes weak convergence on separable spaces and \mathcal{P}_0 is separable [6, p. 72]. Let $d(u, A) \triangleq \inf_{v \in A} d(u, v)$ and $A^\delta \triangleq \{u \in \mathcal{P}_0 | d(u, A) < \delta\}$ so that the Prohorov metric is given by

$$d_P(\mu, \mu') = \inf \{ \delta \in \mathbb{R}_+ | \mu'(A) \leq \mu(A^\delta) + \delta \ \forall \text{ Borel } A \subseteq \mathcal{P}_0 \}.$$

Since the HMP is almost-surely mixing, we can use the fact that $\Pr(d(U_{t+k}, U_t) > C\gamma^t) \leq C\gamma^t$, for all $k \geq 0$, to see that

$$\mu_q^{(t+k)}(A) = \Pr(Q_0 = q, U_{t+k} \in A) \leq \Pr(Q_0 = q, U_t \in A^{C\gamma^t}) + C\gamma^t = \mu_q^{(t)}(A^{C\gamma^t}) + C\gamma^t.$$

This implies that $d_P(\mu_q^{(t)}, \mu_q^{(t+k)}) \leq C\gamma^t$ for all $k \geq 0$. Therefore, $\mu_q^{(t)}(A)$ is a Cauchy sequence with respect to d_P and it converges weakly to some probability measure. Therefore, we can define $\mu_q(A)$ to be the weak limit of $\mu_q^{(t)}(A)$. \square

Definition 2.15. The (forward) *Furstenberg measure* is the unique stationary measure (when it exists) of the joint process $\{Q_t, \alpha_t\}_{t \in \mathbb{Z}}$ and is given by the weak limit

$$\Pr(Q_t = q, \alpha_t \in A) \xrightarrow{w} \mu_q(A),$$

for any Borel measurable set $A \subseteq \mathcal{P}_0$. While this does not depend on the initialization of α_t , one may assume the initialization $\alpha_1 = \pi$ for simplicity.

Remark 2.16. This name is chosen because the measure first appears in the work of Furstenberg and Kifer [15] and is closely related to the work that was started by Furstenberg and Kesten [14].

Consistency of the a posteriori probability (APP)

The following Lemma will be used to make connections between the measures defined in this section.

Lemma 2.17. *Let X, Y be discrete r.v.s and let the APP function be $E_y(x) \triangleq \Pr(X = x | Y = y)$. Then, $E_Y(x) = \Pr(X = x | Y)$ is a random function (due to Y) and we have*

$$\Pr(X = x, E_Y(\cdot) = e(\cdot)) = \Pr(E_Y(\cdot) = e(\cdot)) e(x).$$

Proof. Applying the chain rule and the definition of $E_Y(\cdot)$ gives

$$\begin{aligned} \Pr(X = x, E_Y(\cdot) = e(\cdot)) &= \Pr(E_Y(\cdot) = e(\cdot)) \Pr(X = x | E_Y(\cdot) = e(\cdot)) \\ &= \Pr(E_Y(\cdot) = e(\cdot)) \Pr(X = x | Y) \\ &= \Pr(E_Y(\cdot) = e(\cdot)) e(x), \end{aligned}$$

where the second step follows from the fact that $E_Y(\cdot)$ is a sufficient statistic for X (e.g., X can be faithfully generated from Y using the Markov chain $Y \rightarrow E_Y(\cdot) \rightarrow X$). \square

Proposition 2.18. *The process $\{\alpha_t\}_{t \in \mathbb{Z}}$ forms a Markov chain. If the HMP is almost-surely mixing, then it converges weakly to a unique stationary measure $\mu(A)$.*

Proof. One can see that $\{\alpha_t\}_{t \in \mathbb{Z}}$ is Markov by considering another method of generating the sequence. At each step, we first choose Q_t according to $\alpha_t(\cdot)$, then choose q_{t+1} according to $p_{q_t, q_{t+1}}$, then choose y_t according to $h_{q_t, q_{t+1}}(y_t)$, and finally compute $\alpha_{t+1}(\cdot)$ from $\alpha_t(\cdot)$ and y_t . Of course, this process depends on the initialization of the first α_t but this dependence decays with time if the HMP is almost-surely mixing. For simplicity, one may assume the initialization $\alpha_1 = \pi$ is used.

Comparing this to Proposition 2.14, one see that we are now using $\alpha_t(\cdot)$ as a proxy distribution for Q_t . This works because Lemma 2.17 shows that

$$\Pr(\alpha_t \in A) \inf_{\tilde{\alpha} \in A} \tilde{\alpha}(q) \leq \Pr(Q_t = q, \alpha_t \in A) \leq \Pr(\alpha_t \in A) \sup_{\tilde{\alpha} \in A} \tilde{\alpha}(q),$$

for any open set $A \subseteq \mathcal{P}_0$. By making A arbitrarily small, one can force the LHS and RHS to be arbitrarily close. The proof of weak convergence to a unique stationary distribution as $t \rightarrow \infty$ is essentially identical to the corresponding proof for Proposition 2.14. \square

Definition 2.19. The *(forward) Blackwell measure* is the unique stationary measure (when it exists) of the process $\{\alpha_t\}_{t \in \mathbb{Z}}$ and is given by the weak limit

$$\Pr(\alpha_t \in A) \xrightarrow{w} \mu(A),$$

for any Borel measurable set $A \subseteq \mathcal{P}_0$. From the definition of μ_q , we see also that $\mu(A) = \sum_{q \in \mathcal{Q}} \mu_q(A)$.

Remark 2.20. This name is chosen because this measure first appears in the work of Blackwell [8] and is now commonly called the Blackwell measure [18].

Lemma 2.21. The Radon-Nikodym derivative $\frac{d\mu_q}{d\mu}$ of the (forward) Furstenberg measure μ_q with respect to the (forward) Blackwell measure μ exists and satisfies

$$\frac{d\mu_q}{d\mu}(\alpha) = \Pr(Q_t = q | \alpha_t = \alpha)$$

μ -almost everywhere. This implies that

$$\mu_q(d\alpha) = \alpha(q)\mu(d\alpha).$$

Proof. First, we note that $\mu(A) = \sum_{q \in \mathcal{Q}} \mu_q(A)$ implies that μ_q is absolutely continuous w.r.t. μ . Therefore, the Radon-Nikodym derivative $\frac{d\mu_q}{d\mu}$ exists. Since

$$\frac{\mu_q(A)}{\mu(A)} = \frac{\Pr(Q_t = q, \alpha_t \in A)}{\Pr(\alpha_t \in A)} = \Pr(Q_t = q | \alpha_t \in A),$$

the first result can be seen by choosing A to be arbitrarily small. The second result holds because $\alpha_t(\cdot)$ is the APP estimate of Q_t given $Y_{-\infty}^{t-1}$ and this (e.g., see Lemma 2.17) implies that

$$\Pr(Q_t = q | \alpha_t = \alpha) = \alpha(q).$$

\square

Theorem 2.22 ([8]). In terms of the Blackwell measure, the entropy rate (in nats) of an HMP is

$$H(\mathcal{Y}) = - \int_{\mathcal{P}_0} \mu(d\alpha) \sum_{y \in \mathcal{Y}} \alpha^T M(y) \mathbf{1} \ln(\alpha^T M(y) \mathbf{1}). \quad (2.5)$$

Proof. Consider the sequence $H(Y_t | Y_1^{t-1})$ for any stationary process. This sequence is non-negative and non-increasing and therefore must have a limit. Moreover, the entropy rate

$$H(\mathcal{Y}) \triangleq \lim_{n \rightarrow \infty} \frac{1}{n} H(Y_1, \dots, Y_n) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{t=1}^n H(Y_t | Y_1^{t-1})$$

is the Cesàro mean of this sequence and must have the same limit. Next, we note that

$$\alpha_t^T M(y) \mathbf{1} = \sum_{i,j \in \mathcal{Q}} \alpha_t(i) p_{i,j} h_{i,j}(y) = \Pr(Y_t = y | Y_1^{t-1}).$$

Therefore, (2.5) is simply the expression for $\lim_{t \rightarrow \infty} H(Y_t | Y_1^{t-1})$. \square

Once again, this time in reverse...

One can also reverse time for these Markov processes so that $\{Q_t, \beta_t\}_{t \in \mathbb{Z}}$ forms a backward Markov chain. Starting from q_t and working backwards, one first chooses q_{t-1} according to $\Pr(Q_{t-1} = q_{t-1} | Q_t = q_t) = p_{q_{t-1}, q_t} \pi_{q_{t-1}} / \pi_{q_t}$. Then, one generates y_{t-1} according to $h_{q_{t-1}, q_t}(y_{t-1})$ and computes β_{t-1} from β_t and y_{t-1} .

This process also depends on the initialization of the first β_t but this dependence decays with time if the HMP is almost-surely mixing. For simplicity, one may assume the initialization $\beta_1 = \mathbf{1}$ is used. If the HMP is almost-surely mixing, then the joint distribution of Q_t, β_t converges weakly to a unique stationary distribution as $t \rightarrow -\infty$; the proof is very similar to the corresponding part of the proof of Proposition 2.14. This allows us to define the stationary distribution of the backwards state probability vector.

As with the forward process, we can reduce the state space to $\{\beta_t\}_{t \in \mathbb{Z}}$. At each step, one chooses q_t according to $\Pr(Q_t = q_t) = \beta_t(q_t) \pi_{q_t}$, then continues as described above to generate with q_{t-1}, y_{t-1} , and β_{t-1} . Let $B \subseteq \{u \in \mathcal{A}_0 | \pi^T u = 1\}$ be any open measurable set. Then, using $\beta_t(q) \pi_q$ as a proxy distribution for Q_t works because Lemma 2.17 shows that

$$\Pr(\beta_t \in B) \pi(q) \inf_{\tilde{\beta} \in B} \tilde{\beta}(q) \leq \Pr(Q_t = q, \beta_t \in B) \leq \Pr(\beta_t \in B) \pi(q) \sup_{\tilde{\beta} \in B} \tilde{\beta}(q),$$

and choosing B arbitrarily small allows the LHS and RHS to be made arbitrarily close. This process also depends on the initialization of β_t , but if the HMP is almost-surely mixing, then it converges weakly to a unique stationary distribution.

Definition 2.23. The *backward Furstenberg measure*, is the unique stationary measure (when it exists) of the backwards process $\{Q_t, \beta_t\}_{t \in \mathbb{Z}}$ and is given by the weak limit

$$\Pr(Q_t = q, \beta_t \in B) \xrightarrow{w} \nu_q(B),$$

for any Borel measurable set $B \subseteq \{u \in \mathcal{A}_0 | \pi^T u = 1\}$.

Definition 2.24. The *backward Blackwell measure*, is the unique stationary measure (when it exists) of the backwards process $\{\beta_t\}_{t \in \mathbb{Z}}$ and is given by the weak limit

$$\Pr(\beta_t \in B) \xrightarrow{w} \nu(B),$$

for any Borel measurable set $B \subseteq \{u \in \mathcal{A}_0 | \pi^T u = 1\}$. From the definition of ν_q , we see also that $\nu(B) = \sum_{q \in \mathcal{Q}} \nu_q(B)$.

Lemma 2.25. The Radon-Nikodym derivative $\frac{d\nu_q}{d\nu}$ of the backwards Furstenberg measure ν_q with respect to the backwards Blackwell measure ν exists and satisfies

$$\frac{d\nu_q}{d\nu}(\beta) = \Pr(Q_t = q | \beta_t = \beta)$$

ν -almost everywhere. This implies that

$$\nu_q(d\beta) = \pi(q) \beta(q) \nu(d\beta).$$

Proof. First, we note that $\nu(B) = \sum_{q \in \mathcal{Q}} \nu_q(B)$ implies that ν_q is absolutely continuous w.r.t. ν . Therefore, the Radon-Nikodym derivative $\frac{d\nu_q}{d\nu}$ exists. Since

$$\frac{\nu_q(B)}{\nu(B)} = \frac{\Pr(Q_t = q, \beta_t \in B)}{\Pr(\beta_t \in B)} = \Pr(Q_t = q | \beta_t \in B),$$

the first result can be seen by choosing B to be arbitrarily small. The second result holds because $\beta_t(\cdot)$ is the APP estimate of Q_t given Y_t^∞ and this (e.g., see Lemma 2.17) implies that

$$\Pr(Q_t = q | \beta_t = \beta) = \pi(q)\beta(q).$$

□

3 Taking the Derivative

3.1 The Derivative Shortcut

In this section, we introduce a shortcut often used in the statistical physics community. It was introduced to the author by Measson et al. in [28, 29]. It has also been applied to the problem under consideration by Zuk et al. in [44, 12].

Let $D \subset \mathbb{R}$ be a compact set and $g_n : D^n \rightarrow \mathbb{R}$ be a sequence of functions which essentially depend on a single parameter $\theta \in D$ in n different ways. Abusing notation, we also let $g_n : D \rightarrow \mathbb{R}$ be the same function where this dependency is combined so that $g_n(\theta) = g_n(\theta, \dots, \theta)$. The total derivative of g_n can be written as

$$\frac{d}{d\theta} g_n(\theta) = \sum_{i=1}^n \frac{\partial}{\partial \theta_i} g_n(\theta_1, \dots, \theta_n) \Big|_{(\theta_1, \dots, \theta_n) = (\theta, \dots, \theta)}.$$

This motivates us to define

$$g'_n(\theta_1, \dots, \theta_n) \triangleq \sum_{i=1}^n \frac{\partial}{\partial \theta_i} g_n(\theta_1, \dots, \theta_n).$$

Since the abuse of notation is habit forming, we will also define $g'_n(\theta) \triangleq g'_n(\theta, \dots, \theta)$.

The focus on this paper is the limit of these functions as n goes to infinity, so a few technical details are required. If $g_n(\theta) \rightarrow f(\theta)$ uniformly over $\theta \in D$ and $\lim_{n \rightarrow \infty} g'_n(\theta)$ converges uniformly over $\theta \in D$, then it follows that $f'(\theta) = \lim_{n \rightarrow \infty} g'_n(\theta)$ [4]. One might assume that it is necessary to prove uniform convergence for both of these sequences, but the following standard problem in analysis shows that suffices to consider only the sequence of derivatives.

Lemma 3.1. *Let $g_n : D \rightarrow \mathbb{R}$ be a sequence of functions that are continuously differentiable on a compact set $D \subset \mathbb{R}$. If $g_n(\theta_0)$ converges for some $\theta_0 \in D$ and $g'_n(\theta)$ converges uniformly on D , then the limits*

$$\begin{aligned} f(\theta) &\triangleq \lim_{n \rightarrow \infty} g_n(\theta) \\ f'(\theta) &\triangleq \lim_{n \rightarrow \infty} g'_n(\theta). \end{aligned}$$

both exist and are uniformly continuous on D .

Proof. First, we note that each $g'_n(\theta)$ is uniformly continuous because D is compact. Since $g'_n(\theta)$ converges uniformly, we find that $f'(\theta)$ exists and is uniformly continuous (and hence bounded) on D . Interchanging the limit and integral, based on uniform convergence, implies that

$$\lim_{n \rightarrow \infty} [g_n(\theta) - g_n(\theta_0)] = \lim_{n \rightarrow \infty} \int_{\theta_0}^{\theta} g'_n(x) dx = \int_{\theta_0}^{\theta} \lim_{n \rightarrow \infty} g'_n(x) dx = \int_{\theta_0}^{\theta} f'(x) dx = f(\theta) - f(\theta_0).$$

This implies that $g_n(\theta)$ converges to $f(\theta)$. Finally, we note that $f(\theta)$ is uniformly continuous on D because $f'(\theta)$ exists and is bounded on D . □

3.2 Warmup Example: The Derivative of the Log Spectral Radius

The spectral radius of a real matrix M is defined to be

$$\rho(M) \triangleq \lim_{n \rightarrow \infty} \|M^n\|^{1/n}$$

for any matrix norm. Likewise, the log spectral radius (LSR) of a real matrix M is given by

$$\ln \rho(M) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \|M^n\|,$$

for any matrix norm. Moreover, if M has non-negative entries, then

$$\ln \rho(M) = \lim_{n \rightarrow \infty} \frac{1}{n} \log (u^T M^n v)$$

for any vectors $u, v \in \mathcal{A}_0$.

Let M_θ be a mapping from a compact set $D \subset \mathbb{R}$ to the set of non-negative real matrices. Assume further that M_θ has a unique real eigenvalue λ_1 of maximum modulus (i.e., the 2nd largest eigenvalue λ_2 satisfies $|\lambda_2/\lambda_1| \leq \gamma < 1$) for all $\theta \in D$. Using the shorthand notation $M \triangleq M_{\theta^*}$ for $\theta^* \in D$, we let $a, b \in \mathcal{A}$ be left/right (column) eigenvectors of M with eigenvalue $\rho(M)$; they satisfy $a^T M = \rho(M) a^T$ and $M b = \rho(M) b$. In this case, it is known that the derivative of the LSR is given by

$$\left. \frac{d}{d\theta} \ln \rho(M_\theta) \right|_{\theta=\theta^*} = \frac{a^T M'_{\theta^*} b}{a^T M_{\theta^*} b},$$

where $M' \triangleq M'_{\theta^*}$ is the element-wise derivative defined by $[M'_\theta]_{ij} \triangleq \frac{d}{d\theta} [M_\theta]_{ij}$. Of course, one must assume that M' exists and satisfies $\|M'\| < \infty$.

One can prove this by applying the derivative shortcut to $f(\theta) = \log \rho(M_\theta)$ using

$$g_n(\theta_1, \dots, \theta_n) = \frac{1}{n} \ln \left(u^T \left(\prod_{t=1}^n M_{\theta_t} \right) v \right)$$

for any vectors $u, v \in \mathcal{A}_0$. Based on Lemma 3.1, we focus on $g'_n(\theta)$ by writing

$$\begin{aligned} g'_n(\theta^*) &= \sum_{i=1}^n \frac{\partial}{\partial \theta_i} \frac{1}{n} \ln \left(u^T \left(\prod_{t=1}^n M_{\theta_t} \right) v \right) \Big|_{(\theta_1, \dots, \theta_n) = (\theta^*, \dots, \theta^*)} \\ &= \frac{1}{n} \sum_{i=1}^n \frac{\partial}{\partial \theta_i} \ln \left(u^T \left(\prod_{t=1}^{i-1} M_{\theta_t} \right) M(\theta_i) \left(\prod_{t=i+1}^n M_{\theta_t} \right) v \right) \Big|_{\theta_1^n = (\theta^*, \dots, \theta^*)} \\ &= \frac{1}{n} \sum_{i=1}^n \frac{u^T \left(\prod_{t=1}^{i-1} M_{\theta_t} \right) M'_{\theta_i} \left(\prod_{t=i+1}^n M_{\theta_t} \right) v}{u^T \left(\prod_{t=1}^{i-1} M_{\theta_t} \right) M_{\theta_i} \left(\prod_{t=i+1}^n M_{\theta_t} \right) v} \Big|_{\theta_1^n = (\theta^*, \dots, \theta^*)} \\ &= \frac{1}{n} \sum_{i=1}^n \frac{u^T M^{i-1} M' M^{n-i} v}{u^T M^{i-1} M M^{n-i} v}, \end{aligned}$$

where we have used that

$$\frac{d}{d\theta} x^T M_\theta y = \sum_{k,l} x_k \frac{d}{d\theta} [M_\theta]_{k,l} y_l = x^T M'_\theta y.$$

Since M_θ satisfies $|\lambda_2/\lambda_1| \leq \gamma$ for all $\theta \in D$, it follows that

$$\begin{aligned} \frac{u^T M^{i-1}}{\|u^T M^{i-1}\|} &= a^T + O(\gamma^{i-1}) \\ \frac{M^{n-i} v}{\|M^{n-i} v\|} &= b + O(\gamma^{n-i}). \end{aligned}$$

Treating the boundary and interior terms, in the sum, separately gives

$$g'_n(\theta^*) = O\left(\frac{\lfloor (\ln n)^2 \rfloor}{n} \frac{\|M'\| (a^T b)}{\rho(M) \frac{u^T b}{\|u\|} \frac{a^T v}{\|v\|}}\right) + \frac{1}{n} \sum_{i=\lfloor (\ln n)^2 \rfloor + 1}^{n - \lfloor (\ln n)^2 \rfloor} \frac{a^T M' b + O(\gamma^{(\ln n)^2}) \|M'\|}{a^T M b + O(\gamma^{(\ln n)^2}) \|M\|}.$$

Therefore, $g_n(\theta)$ and $g'_n(\theta)$ converge uniformly for all $\theta \in D$ and we find that

$$f'(\theta^*) = \frac{a^T M' b}{a^T M b}.$$

3.3 The Derivative of the Entropy Rate

Let $M_\theta(y)$ be transition observation probability matrix of an HMP, which depends on the real parameter θ , and let π be the stationary distribution of the underlying Markov chain. To compute the derivative of the entropy rate, we define

$$\begin{aligned} g_n(\theta_1, \dots, \theta_n) &= -\frac{1}{n} \sum_{y_1^n \in \mathcal{Y}^n} \Pr(Y_1^n = y_1^n; \theta_1^n) \ln \Pr(Y_1^n = y_1^n; \theta_1^n) \\ &= -\frac{1}{n} \sum_{y_1^n \in \mathcal{Y}^n} \pi^T \left(\prod_{i=1}^n M_{\theta_i}(y_i) \right) \mathbf{1} \cdot \ln \left[\pi^T \left(\prod_{i=1}^n M_{\theta_i}(y_i) \right) \mathbf{1} \right] \Big|_{(\theta_1, \dots, \theta_n) = (\theta^*, \dots, \theta^*)}. \end{aligned}$$

This implies that $f(\theta) = \lim_{n \rightarrow \infty} g_n(\theta) = H(\mathcal{Y}; \theta)$ in nats.

Theorem 3.2. *Let $D \subset \mathbb{R}$ be a compact set and assume that $\frac{d}{d\theta} \pi = \mathbf{0}$ and $M'_\theta(y) \triangleq \frac{d}{d\theta} M_\theta(y)$ exists for all $\theta \in D$. Then, if the HMP is well-defined and ϵ -primitive for all $\theta \in D$, then $f'(\theta^*) = \frac{d}{d\theta} H(\mathcal{Y}; \theta) \Big|_{\theta=\theta^*}$ equals*

$$- \int_{\mathcal{A}_0} \mu(d\alpha) \int_{\mathcal{A}_0} \nu(d\beta) \sum_{y \in \mathcal{Y}} \alpha^T M'_{\theta^*}(y) \beta \ln(\alpha^T M_{\theta^*}(y) \beta), \quad (3.1)$$

where μ and ν are the forward/backward Blackwell measures of the HMP at $\theta = \theta^*$. Moreover, $f(\theta)$ and $f'(\theta)$ are uniformly continuous on D .

Proof. The following shorthand is used throughout: $\pi_t(q) \triangleq \Pr(Q_t = q)$, $M(y) \triangleq M_{\theta^*}(y)$, $M'(y) \triangleq M'_{\theta^*}(y)$, and $M(y_j^k) \triangleq \prod_{t=j}^k M_{\theta^*}(y_t)$. For the HMP to be well-defined, the transition matrices must satisfy $\sum_{y \in \mathcal{Y}} M_\theta(y) \mathbf{1} = \mathbf{1}$ and $\sum_{y \in \mathcal{Y}} M'_\theta(y) \mathbf{1} = \mathbf{0}$ for all $\theta \in D$. It follows that, for any $u \in \mathcal{P}_0$, one has

$$\begin{aligned} \sum_{y_1^n \in \mathcal{Y}^n} u^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} &= 1 \\ \frac{\partial}{\partial \theta_j} \sum_{y_1^n \in \mathcal{Y}^n} u^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} &= 0. \end{aligned} \quad (3.2)$$

Based on Lemma 3.1, we note that the entropy rate exists for all $\theta \in D$ and focus on the derivative

$$\begin{aligned} g'_n(\theta^*) &\stackrel{(a)}{=} -\frac{1}{n} \sum_{j=1}^n \frac{\partial}{\partial \theta_j} \sum_{y_1^n \in \mathcal{Y}^n} \pi_1^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \cdot \left(\ln \left[C_j \pi_1^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \right] - \ln C_j \right) \Big|_{\theta_j = \theta^*} \\ &\stackrel{(b)}{=} -\frac{1}{n} \sum_{j=1}^n \frac{\partial}{\partial \theta_j} \sum_{y_1^n \in \mathcal{Y}^n} \pi_1^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \cdot \ln \left[C_j \pi_1^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \right] \Big|_{\theta_j = \theta^*} \\ &\stackrel{(c)}{=} -\frac{1}{n} \sum_{j=1}^n \frac{\partial}{\partial \theta_j} \sum_{y_1^n \in \mathcal{Y}^n} \pi_1^T M(y_1^{j-1}) M_{\theta_j}(y_j) M(y_{j+1}^n) \mathbf{1} \cdot \ln \left[\frac{\pi_1^T M(y_1^{j-1}) M_{\theta_j}(y_j) M(y_{j+1}^n) \mathbf{1}}{(\pi_1^T M(y_1^{j-1}) \mathbf{1}) (\pi_{j+1}^T M(y_{j+1}^n) \mathbf{1})} \right] \Big|_{\theta_j = \theta^*}, \end{aligned}$$

where (a) holds for arbitrary positive values C_1, \dots, C_n , (b) follows because (3.2) implies the $\ln C_j$ gives no contribution if $\frac{\partial}{\partial \theta_j} C_j = 0$, and (c) follows from choosing

$$C_j = \Pr(Y_1^{j-1} = y_1^{j-1}) \Pr(Y_{j+1}^n = y_{j+1}^n) = \left(\pi_1^T M(y_1^{j-1}) \mathbf{1} \right) \left(\pi_{j+1}^T M(y_{j+1}^n) \mathbf{1} \right).$$

One subtlety is that $\pi_{j+1} = \pi_j \sum_{y \in \mathcal{Y}} M_{\theta_j}(y)$ is affected by θ_j . So, small changes in θ_j cause small changes in π_{j+1} and we must add the condition $\frac{d}{d\theta} \pi = \mathbf{0}$ to guarantee that $\frac{\partial}{\partial \theta_j} C_j = 0$. After adding this condition, we may safely assume that $\pi_j = \pi$ for $j = 1, \dots, n$. See Remark 3.3 for more details.

For Borel measurable sets $A \subseteq \mathcal{A}_0$ and $B \subseteq \{u \in \mathcal{A}_0 \mid \pi^T u = 1\}$, the sets

$$U_j(A) \triangleq \left\{ y_1^{j-1} \in \mathcal{Y}^{j-1} \mid \alpha_j^T = \frac{\pi^T M(y_1^{j-1})}{\pi^T M(y_1^{j-1}) \mathbf{1}} \in A \right\}$$

$$V_j(B) \triangleq \left\{ y_j^n \in \mathcal{Y}^{n-j+1} \mid \beta_j = \frac{M(y_j^n) \mathbf{1}}{\pi^T M(y_j^n) \mathbf{1}} \in B \right\}$$

will be used to define the measures $\mu^{(j)}(A) \triangleq \Pr(Y_1^{j-1} \in U_j(A))$ and $\nu^{(j)}(B) \triangleq \Pr(Y_j^n \in V_j(B))$ for the forward/backward state probabilities. In this case, $\mu^{(j)}(\cdot), \nu^{(j)}(\cdot)$ are probability measures on \mathcal{A}_0 for the random variables α_j, β_j . Using these measures, we find that $g'_n(\theta^*)$ is given by

$$\begin{aligned} &= -\frac{1}{n} \sum_{j=1}^n \frac{\partial}{\partial \theta_j} \sum_{y_1^n \in \mathcal{Y}^n} \overbrace{\pi^T M(y_1^{j-1})}^{\alpha_j^T \cdot \pi^T M(y_1^{j-1}) \mathbf{1}} M_{\theta_j}(y_j) \overbrace{M(y_{j+1}^n) \mathbf{1}}^{\beta_{j+1} \cdot \pi^T M(y_{j+1}^n) \mathbf{1}} \ln \left[\frac{\overbrace{\pi^T M(y_1^{j-1})}^{\alpha_j^T}}{\pi^T M(y_1^{j-1}) \mathbf{1}} M_{\theta_j}(y_j) \frac{\overbrace{M(y_{j+1}^n) \mathbf{1}}^{\beta_{j+1}}}{\pi^T M(y_{j+1}^n) \mathbf{1}} \right] \Bigg|_{\theta_j = \theta^*} \\ &= -\frac{1}{n} \sum_{j=1}^n \int_{\mathcal{A}_0} \mu^{(j)}(d\alpha) \int_{\mathcal{A}_0} \nu^{(j+1)}(d\beta) \frac{\partial}{\partial \theta_j} \sum_{y_j \in \mathcal{Y}} \alpha^T M_{\theta_j}(y_j) \beta \ln(\alpha^T M_{\theta_j}(y_j) \beta) \Bigg|_{\theta_j = \theta^*} \\ &= -\frac{1}{n} \sum_{j=1}^n \int_{\mathcal{A}_0} \mu^{(j)}(d\alpha) \int_{\mathcal{A}_0} \nu^{(j+1)}(d\beta) \sum_{y_j \in \mathcal{Y}} [\alpha^T M'(y_j) \beta \ln(\alpha^T M(y_j) \beta) + \alpha^T M'(y_j) \beta]. \end{aligned}$$

All that is left is to compute the sum. If the HMP is almost-surely mixing, then the results of Section 2.4 show that measures converge weakly (i.e., $\mu^{(j)} \rightarrow \mu$ and $\nu^{(j)} \rightarrow \nu$). Moreover, Lemma A.2 in Appendix A.1 shows that the convergence rate is exponential. Therefore, most of the terms in the sum have essentially the same value. Like the LSR, we neglect terms within $(\ln n)^2$ of the block edge because their contribution is negligible as $n \rightarrow \infty$. The exponential convergence of the stationary measures also shows that the interior terms become equal at the super polynomial rate $\gamma^{(\ln n)^2} = n^{\ln n \cdot \ln \gamma}$. Therefore, $f_n(\theta)$ and $f'_n(\theta)$ converge uniformly for all $\theta \in D$ and

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \sum_{j=1}^n \int_{\mathcal{A}_0} \mu^{(j)}(d\alpha) \int_{\mathcal{A}_0} \nu^{(j+1)}(d\beta) \sum_{y_j \in \mathcal{Y}} [\alpha^T M'(y_j) \beta \ln(\alpha^T M(y_j) \beta) + \alpha^T M'(y_j) \beta]$$

converges to

$$\frac{d}{d\theta} H(\mathcal{Y}; \theta) \Big|_{\theta = \theta^*} = - \int_{\mathcal{A}_0} \mu(d\alpha) \int_{\mathcal{A}_0} \nu(d\beta) \sum_{y \in \mathcal{Y}} [\alpha^T M'(y) \beta \ln(\alpha^T M(y) \beta) + \alpha^T M'(y) \beta]. \quad (3.3)$$

Finally, the last term in (3.3) is shown to be zero in Lemma 3.4. \square

Remark 3.3. The necessity of the condition $\frac{d}{d\theta} \pi = \mathbf{0}$ in Theorem 3.2 can be a bit subtle. This is because the π -term in many equations (e.g., $\pi^T M(y_{j+1}^n) \mathbf{1}$) actually represents the state distribution at a

particular time (e.g., time $j+1$). The indices are dropped after the first few steps because the underlying Markov chain is stationary and the state distribution is independent of time. For example, the proof liberally uses the assumption that

$$\Pr(Y_{j+1}^n = y_{j+1}^n) = \sum_{q, q' \in Q} \Pr(Q_{j+1} = q) \Pr(Q_{n+1} = q', Y_{j+1}^n = y_{j+1}^n | Q_{j+1} = q) = \pi M(y_{j+1}^n) \mathbf{1},$$

where the last step clearly requires that $\Pr(Q_{j+1} = q) = \pi(q)$. Moreover, this is not simply a problem with the proof. The author has applied the formula from Theorem 3.2 to a Markov chain (where the true entropy-rate derivative is well-known) and shown that the two expressions become equal only if $\frac{d}{d\theta}\pi = \mathbf{0}$.

Lemma 3.4. *The following properties of the forward/backward Blackwell measures will be useful:*

$$\begin{aligned} \int_{\mathcal{A}_0} \mu(d\alpha) \alpha &= \pi \\ \int_{\mathcal{A}_0} \nu(d\beta) \beta &= \mathbf{1} \\ \int_{\mathcal{A}_0} \mu(d\alpha) \sum_{y \in \mathcal{Y}} \alpha^T M(y) \beta &= 1 \\ \int_{\mathcal{A}_0} \nu(d\beta) \sum_{y \in \mathcal{Y}} \alpha^T M(y) \beta &= 1 \\ \int_{\mathcal{A}_0} \mu(d\alpha) \int_{\mathcal{A}_0} \nu(d\beta) \sum_{y \in \mathcal{Y}} \alpha^T M'(y) \beta &= 0 \end{aligned}$$

Proof. The proof is deferred to the appendix. □

3.4 Behavior of the Entropy Rate in the High Noise Regime

Suppose the domain of θ includes a “high noise” point θ^* where the channel output provides no information about the channel state. In this case, the forward/backward Blackwell measures become singletons on $\pi, \mathbf{1}$ and the entropy rate $H(\mathcal{Y}; \theta)$ converges to the single-letter entropy $H(Y; \theta)$ as $\theta \rightarrow \theta^*$. In the high-noise regime, one can also evaluate the derivative from Theorem 3.2 in closed form and extend the formula to the 2nd derivative. In this section, we compare the expansions of $H(\mathcal{Y}; \theta)$ and $H(Y; \theta)$.

First, we consider the single-letter entropy

$$\begin{aligned} H(Y; \theta) &= - \sum_{y \in \mathcal{Y}} \Pr(Y_t = y) \log(\Pr(Y_t = y)) \\ &= - \sum_{y \in \mathcal{Y}} \pi^T M_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}), \end{aligned}$$

where π is the stationary distribution of the underlying Markov chain as a function of θ .

Lemma 3.5. *Under the assumption that $\frac{d}{d\theta}\pi = \mathbf{0}$ for all $\theta \in D$, the 1st derivative w.r.t. θ of the single-letter entropy is given by*

$$\frac{d}{d\theta} H(Y; \theta) = - \sum_{y \in \mathcal{Y}} \pi^T M'_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}),$$

Under the same assumption, the 2nd derivative w.r.t. θ is given by

$$\frac{d^2}{d\theta^2} H(Y; \theta) = - \sum_{y \in \mathcal{Y}} \frac{(\pi^T M'_\theta(y) \mathbf{1})^2}{\pi^T M_\theta(y) \mathbf{1}} - \sum_{y \in \mathcal{Y}} \pi^T M''_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}). \quad (3.4)$$

Proof. In particular, the 1st derivative is given by

$$\begin{aligned}\frac{d}{d\theta}H(Y; \theta) &= -\frac{d}{d\theta} \sum_{y \in \mathcal{Y}} \pi^T M_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}) \\ &= -\sum_{y \in \mathcal{Y}} \left(\left(\frac{d}{d\theta} \pi^T \right) M_\theta(y) \mathbf{1} + \pi^T M'_\theta(y) \mathbf{1} \right) \log(\pi^T M_\theta(y) \mathbf{1}) - \frac{d}{d\theta} \sum_{y \in \mathcal{Y}} \pi^T M_\theta(y) \mathbf{1} \\ &= -\sum_{y \in \mathcal{Y}} \pi^T M'_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1})\end{aligned}$$

because $\frac{d}{d\theta} \pi = \mathbf{0}$ and $\sum_{y \in \mathcal{Y}} \pi^T M_\theta(y) \mathbf{1} = 1$ for all θ . Since $\frac{d}{d\theta} \pi = \mathbf{0}$ for all $\theta \in D$, the 2nd derivative is given by

$$\begin{aligned}\frac{d^2}{d\theta^2}H(Y; \theta) &= -\frac{d}{d\theta} \sum_{y \in \mathcal{Y}} \pi^T M'_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}) \\ &= -\sum_{y \in \mathcal{Y}} \pi^T M''_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}) - \sum_{y \in \mathcal{Y}} \frac{(\pi^T M'_\theta(y) \mathbf{1})^2}{\pi^T M_\theta(y) \mathbf{1}}.\end{aligned}$$

□

Now, we consider closed form evaluation of Theorem 3.2. Since the first derivative is often zero at $\theta = \theta^*$, we are fortunate that a new formula for the 2nd derivative can also be evaluated in closed form.

Theorem 3.6. *If there is a function $s(y)$, a $\theta^* \in D$, and a matrix P such that $\lim_{\theta \rightarrow \theta^*} M(y) = s(y)P$ for all $y \in \mathcal{Y}$, then*

$$\frac{d}{d\theta}H(\mathcal{Y}; \theta) \Big|_{\theta=\theta^*} = -\sum_{y \in \mathcal{Y}} \pi^T M'(y) \mathbf{1} \ln(s(y))$$

and

$$\frac{d^2}{d\theta^2}H(\mathcal{Y}; \theta) \Big|_{\theta=\theta^*} = -\sum_{y \in \mathcal{Y}} \pi^T M''(y) \mathbf{1} \ln(s(y)) - \sum_{y \in \mathcal{Y}} \frac{(\pi^T M'(y) \mathbf{1})^2}{\pi^T M(y) \mathbf{1}}. \quad (3.5)$$

Proof. The proof is deferred to the appendix. □

3.5 HMP Example: A Binary Markov-1 Source with BSC Noise

Consider the HMP defined by a binary Markov-1 source observed through a BSC(ε). The two-state Markov process is defined by $\Pr(Q_{t+1} = j | Q_t = i) = p_{ij}$ with stationary distribution $\Pr(Q_t = i) = \pi(i)$, and $\pi(0) = 1 - \pi(1) = \frac{1-p_{11}}{2-p_{00}-p_{11}}$. The output of the HMP is simply the observation of state through a BSC or more specifically

$$h_{i,j}(y) = \begin{cases} 1 - \varepsilon & \text{if } y = i \\ \varepsilon & \text{otherwise} \end{cases}.$$

The entropy rate of this process was considered earlier using a range of techniques [31, 32, 18, 44]. Now, we will consider the entropy rate of this process as $\varepsilon \rightarrow \frac{1}{2}$ (i.e., in the high-noise regime). This special case was also treated earlier and very similar results were obtained using different methods in [20, 19, 33].

Since we are interested in the high-noise regime, we start by analyzing the system using the upper bound $H(\mathcal{Y}) \leq H(Y)$. This gives

$$H(Y) = -\sum_{y \in \mathcal{Y}} \Pr(Y = y) \ln(\Pr(Y = y)),$$

where

$$\begin{aligned}\Pr(Y = 0) &= \pi(0)p_{00}(1 - \varepsilon) + \pi(0)p_{01}\varepsilon + \pi(1)p_{10}(1 - \varepsilon) + \pi(1)p_{11}\varepsilon \\ \Pr(Y = 1) &= \pi(0)p_{00}\varepsilon + \pi(0)p_{01}(1 - \varepsilon) + \pi(1)p_{10}\varepsilon + \pi(1)p_{11}(1 - \varepsilon).\end{aligned}$$

Using the Taylor expansion of $H(Y; \theta)$ around $\theta = \frac{1}{2} - \varepsilon$, we find that

$$H(\mathcal{Y}) \leq H(Y; \theta) = \ln 2 - \frac{4(p_{00}^2 - p_{11}^2)}{(2 - p_{00} - p_{11})^2} \frac{\theta^2}{2} + O(\theta^4). \quad (3.6)$$

To calculate this expansion exactly for $H(\mathcal{Y})$, we apply Theorem 3.6. The conditions of the Theorem are satisfied because

$$\begin{aligned} M_\theta(y) &= \begin{cases} \begin{bmatrix} p_{00}(1-\varepsilon) & p_{01}\varepsilon \\ p_{10}(1-\varepsilon) & p_{11}\varepsilon \end{bmatrix} & \text{if } y = 0 \\ \begin{bmatrix} p_{00}\varepsilon & p_{01}(1-\varepsilon) \\ p_{10}\varepsilon & p_{11}(1-\varepsilon) \end{bmatrix} & \text{if } y = 1 \end{cases} \\ &= \begin{cases} \frac{1}{2} \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} + \theta \begin{bmatrix} p_{00} & -p_{01} \\ p_{10} & -p_{11} \end{bmatrix} & \text{if } y = 0 \\ \frac{1}{2} \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} - \theta \begin{bmatrix} p_{00} & -p_{01} \\ p_{10} & -p_{11} \end{bmatrix} & \text{if } y = 1 \end{cases} \end{aligned}$$

implies $M_\theta(0) = M_\theta(1)$ at $\theta = 0$ (i.e., $\varepsilon = \frac{1}{2}$). Computing (3.5), which is simplified by the symmetry of $M_\theta(y)$ and the fact that $M_\theta''(y)$ is the zero matrix, gives

$$\begin{aligned} \left. \frac{d^2}{d\theta^2} H(\mathcal{Y}; \theta) \right|_{\theta=0} &= -2 \frac{\left(\begin{bmatrix} \frac{1-p_{11}}{2-p_{00}-p_{11}} & \frac{1-p_{00}}{2-p_{00}-p_{11}} \end{bmatrix} \begin{bmatrix} p_{00} & -p_{01} \\ p_{10} & -p_{11} \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right)^2}{\frac{1}{2} \begin{bmatrix} \frac{1-p_{11}}{2-p_{00}-p_{11}} & \frac{1-p_{00}}{2-p_{00}-p_{11}} \end{bmatrix} \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix}} \\ &= -\frac{4(p_{00}^2 - p_{11}^2)}{(2 - p_{00} - p_{11})^2}. \end{aligned} \quad (3.7)$$

Since $H(\mathcal{Y}; 0) = \ln 2$, this implies that the upper bound is tight with respect to the first non-zero term in the high-noise expansion.

3.6 Example 2: A Conditionally Gaussian HMP

Consider an HMP where the output distribution, conditioned on the state of underlying Markov chain, is Gaussian. Suppose that the Gaussian associated with the transition from state i to state j has mean $\theta \cdot m_{ij}$ and variance 1, then this implies that $h_{ij}(y) = \frac{1}{\sqrt{2\pi}} e^{-(y-\theta m_{ij})^2/2}$. Since the HMP loses state dependence as $\theta \rightarrow 0$, we first consider the derivatives w.r.t. θ of the single-letter entropy

$$H(Y; \theta) = - \int_{-\infty}^{\infty} \pi^T M_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}) dy.$$

In this case, the stationary distribution does not depend on θ so translating Lemma 3.5 to the continuous alphabet case gives

$$\begin{aligned} \left. \frac{d}{d\theta} H(Y; \theta) \right|_{\theta=0} &= - \lim_{\theta \rightarrow 0} \int_{-\infty}^{\infty} \pi^T M'_\theta(y) \mathbf{1} \log(\pi^T M_\theta(y) \mathbf{1}) dy \\ &= - \lim_{\theta \rightarrow 0} \int_{-\infty}^{\infty} \sum_{i,j \in \mathcal{Q}} \frac{\pi(i)p_{ij}}{\sqrt{2\pi}} e^{-(y-\theta m_{ij})^2/2} m_{ij}(y - \theta m_{ij}) \log \left(\sum_{k,l \in \mathcal{Q}} \frac{\pi(k)p_{kl}}{\sqrt{2\pi}} e^{-(y-\theta m_{kl})^2/2} \right) dy \\ &= - \int_{-\infty}^{\infty} \sum_{i,j \in \mathcal{Q}} \pi(i)p_{ij} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} m_{ij} y \log \left(\frac{1}{\sqrt{2\pi}} e^{-y^2/2} \right) dy \\ &= - \int_{-\infty}^{\infty} \sum_{i,j \in \mathcal{Q}} \pi(i)p_{ij} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} m_{ij} \left[y \log \left(\frac{1}{\sqrt{2\pi}} \right) - \frac{y^3}{2} \right] dy \\ &= 0, \end{aligned}$$

because the odd moments of a zero-mean Gaussian are zero. Likewise, the formula for 2nd derivative (3.4) can be translated into

$$\frac{d^2}{d\theta^2}H(Y; \theta) = - \int_{-\infty}^{\infty} \pi^T M_{\theta}''(y) \mathbf{1} \log(\pi^T M_{\theta}(y) \mathbf{1}) dy - \int_{-\infty}^{\infty} \frac{(\pi^T M_{\theta}'(y) \mathbf{1})^2}{\pi^T M_{\theta}(y) \mathbf{1}} dy$$

The second term T_2 of the expression for $\frac{d^2}{d\theta^2}H(Y; \theta)|_{\theta=0}$ is given by

$$\begin{aligned} T_2 &= - \lim_{\theta \rightarrow 0} \int_{-\infty}^{\infty} \frac{\left(\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} \frac{1}{\sqrt{2\pi}} e^{-(y-\theta m_{ij})^2/2} m_{ij} (y - \theta m_{ij}) \right)^2}{\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} \frac{1}{\sqrt{2\pi}} e^{-(y-\theta m_{ij})^2/2}} dy \\ &= - \int_{-\infty}^{\infty} \frac{\left(\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} m_{ij} y \right)^2}{\frac{1}{\sqrt{2\pi}} e^{-y^2/2}} dy \\ &= - \left(\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij} \right)^2 \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} y^2 dy \\ &= - \left(\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij} \right)^2. \end{aligned}$$

Using the fact that

$$\pi^T M_{\theta}''(y) \mathbf{1} = \sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} \frac{1}{\sqrt{2\pi}} e^{-(y-\theta m_{ij})^2/2} m_{ij}^2 [(y - \theta m_{ij})^2 - 1],$$

we can write the first term T_1 of the expression for $\frac{d^2}{d\theta^2}H(Y; \theta)|_{\theta=0}$ as

$$\begin{aligned} T_1 &= - \lim_{\theta \rightarrow 0} \int_{-\infty}^{\infty} \pi^T M_{\theta}''(y) \mathbf{1} \log(\pi^T M_{\theta}(y) \mathbf{1}) dy \\ &= - \int_{-\infty}^{\infty} \sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} m_{ij}^2 (y^2 - 1) \log\left(\frac{1}{\sqrt{2\pi}} e^{-y^2/2}\right) dy \\ &\stackrel{(a)}{=} \frac{1}{2} \sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij}^2 \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} (y^4 - y^2) dy \\ &= \sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij}^2, \end{aligned}$$

where (a) follows from the fact that the 4th moment of a standard Gaussian is 3.

Comparing Lemma 3.5 with Theorem 3.6 shows that the first two terms in the expansion of $H(Y; \theta)$ match the first two terms in the expansion of $H(\mathcal{Y}; \theta)$ at $\theta = 0$. Therefore, we have

$$\left. \frac{d^2}{d\theta^2} H(\mathcal{Y}; \theta) \right|_{\theta=0} = \left. \frac{d^2}{d\theta^2} H(Y; \theta) \right|_{\theta=0} = \sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij}^2 - \left(\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij} \right)^2. \quad (3.8)$$

4 Application: High-Noise Capacity Expansions for FSCs

4.1 The Derivative of Capacity for an FSC

Now, we will use the previous result to compute the derivative of the capacity. The mutual information $I(X; Y)$ between the r.v.s X and Y is defined by $I(X; Y) \triangleq H(Y) - H(Y|X)$, where the conditional

entropy is defined by $H(Y|X) \triangleq H(X, Y) - H(X)$. Since the mutual information depends on the input distribution, the capacity is defined to be the supremum of the mutual information over all input distributions [11]. Therefore, some care must be taken when expressing the derivative of the capacity in terms of the derivative of the mutual information.

Consider a family of FSCs whose entropy rate is differentiable with respect to some parameter θ . Let the input distribution be Markov with memory m (e.g., defined by the vector \vec{P} containing $|\mathcal{X}|^{m+1}$ values) and the optimal input distribution be $\vec{P}(\theta)$. In this case, we let the mutual information rate be $\mathcal{I}(\theta, \vec{P})$ and the Markov- m capacity be $\mathcal{C}(\theta) = \mathcal{I}(\theta, \vec{P}(\theta))$.

Lemma 4.1. *The derivative of the Markov- m capacity is given by*

$$\frac{d}{d\theta}\mathcal{C}(\theta) = \frac{d}{d\theta}\mathcal{I}(\theta, \vec{P}(\theta)) = \mathcal{I}'_{\theta}(\theta, \vec{P}(\theta)), \quad (4.1)$$

where $\mathcal{I}'_{\theta}(\theta, \vec{P}(\theta))$ is the derivative (w.r.t. θ) of the mutual information rate evaluated at the capacity achieving input distribution for θ .

Proof. Expanding the derivative of $\mathcal{C}(\theta)$ in terms of $\mathcal{I}'_{\theta}(\theta, \vec{P})$ and the gradient vector $\mathcal{I}'_{\vec{P}}(\theta, \vec{P})$ (w.r.t. input distribution), gives

$$d\mathcal{I}(\theta, \vec{P}) = \mathcal{I}'_{\theta}(\theta, \vec{P}) d\theta + \mathcal{I}'_{\vec{P}}(\theta, \vec{P}) \cdot d\vec{P}.$$

The optimality of $\vec{P}(\theta)$ implies $\mathcal{I}'_{\vec{P}}(\theta, \vec{P}(\theta)) \cdot d\vec{P} = 0$ for any $d\vec{P}$ satisfying $d\vec{P} \cdot \mathbf{1} = 0$ (i.e., the sum of $\vec{P}(\theta)$ is a constant). So, the derivative of the capacity is the derivative of the mutual information rate and we have (4.1). \square

Corollary 4.2. *If there is a “high noise” point $\theta^* \in D$ where the Markov- m capacity satisfies $\mathcal{C}(\theta^*) = 0$ and $\mathcal{C}'(\theta^*) = 0$, then*

$$\left. \frac{d^2}{d\theta^2}\mathcal{C}(\theta) \right|_{\theta=\theta^*} = \mathcal{I}''_{\theta}(\theta^*, \vec{P}(\theta^*)),$$

where $\mathcal{I}''_{\theta}(\theta^*, \vec{P}(\theta^*))$ is the 2nd derivative (w.r.t. θ) of the mutual information rate evaluated at the capacity achieving input distribution for θ .

Proof. First, we write the 2nd derivative as

$$\begin{aligned} \left. \frac{d^2}{d\theta^2}\mathcal{C}(\theta) \right|_{\theta=\theta^*} &= \lim_{\theta \rightarrow \theta^*} \frac{d}{d\theta} \mathcal{I}'_{\theta}(\theta, \vec{P}(\theta)) \\ &= \mathcal{I}''_{\theta}(\theta^*, \vec{P}(\theta^*)) + \lim_{\theta \rightarrow \theta^*} \left[\frac{d}{d\vec{P}} \mathcal{I}'_{\theta}(\theta, \vec{P}) \right]_{\vec{P}=\vec{P}(\theta^*)} \cdot \vec{P}'(\theta^*). \end{aligned}$$

Now, recall that $\mathcal{I}'_{\theta}(\theta^*, \vec{P}(\theta^*)) = 0$ and suppose that the 2nd term is positive. In this case, a small change in \vec{P} in the direction $\vec{P}'(\theta^*)$ must give an $\mathcal{I}'_{\theta}(\theta^*, \vec{P}) > 0$. But, this contradicts the fact that

$$0 = \mathcal{C}'(\theta^*) \geq \max_{\vec{P}} \mathcal{I}'_{\theta}(\theta^*, \vec{P}).$$

Therefore, the 2nd term must be zero. \square

If the domain of θ includes a “high noise” point θ^* where the channel output provides no information about the channel state, then Theorem 3.6 shows that the first two θ -derivatives of the entropy rate $H(\mathcal{Y}; \theta)$ can be calculated at $\theta = \theta^*$. In fact, one also sees that they match the first two θ -derivatives of the single-letter entropy $H(Y; \theta)$ at $\theta = \theta^*$. Using Lemma 4.1 and Corollary 4.2, we see that these derivatives also equal the derivative of the Markov- m capacity in this case. But this equality holds for all m , so we can take a limit to see that it must hold also for the true capacity [10]. Even without this, however, we can use the fact that $H(\mathcal{Y}; \theta) \leq H(Y; \theta)$ to upper bound the maximum entropy rate over all input distributions.

4.2 FSC Example: A BSC with an RLL Constraint

Consider the FSC defined by the BSC(ε) with a (0,1) run-length (RLL) constraint [24]. This is a standard binary symmetric channel with a constraint that the input cannot have two 1s in a row (e.g., this requires a two-state input process). The two-state input process is defined by $\Pr(X_{t+1} = j | X_t = i) = p_{ij}$ with $p_{11} = 0$, $\Pr(X_t = i) = \pi(i)$, and $\pi(0) = 1 - \pi(1) = \frac{1}{2-p_{00}}$.

The mutual information rate between the input and output satisfies

$$\begin{aligned} I(\mathcal{X}; \mathcal{Y}) &= H(\mathcal{Y}) - H(\mathcal{Y}|\mathcal{X}) \\ &\leq H(Y_i) - h(\varepsilon), \end{aligned}$$

where $h(\varepsilon) = -\varepsilon \ln \varepsilon - (1 - \varepsilon) \ln(1 - \varepsilon)$ is the binary entropy function in nats. Now, we can let $\theta = \frac{1}{2} - \varepsilon$ and combine the entropy-rate expansion from (3.6) with the fact that $h(\frac{1}{2} - \theta) = \ln 2 - 2\theta^2 + O(\theta^4)$. The resulting high-noise expansion for the upper bound is

$$I(\mathcal{X}; \mathcal{Y}) \leq \frac{8(1-p_{00})}{(2-p_{00})^2} \theta^2 + O(\theta^4).$$

Notice that the leading coefficient achieves a unique maximum value of $\frac{2}{\ln 2}$ at $p_{00} = 0$. Since this upper bound only depends on the single-letter probabilities, it cannot be increased by extending the memory of the input process.

To see that this rate is achievable, we apply Theorem 3.6 to our system. Taking the result from (3.7), we find that

$$\begin{aligned} I(\mathcal{X}; \mathcal{Y}) &= H(\mathcal{Y}) - H(\mathcal{Y}|\mathcal{X}) \\ &= \left(1 - \frac{2p_{00}^2}{(2-p_{00})^2} \theta^2 + o(\theta^2) \right) - (1 - 2\theta^2 + O(\theta^4)) \\ &= \frac{8(1-p_{00})}{(2-p_{00})^2} \theta^2 + o(\theta^2). \end{aligned}$$

So the leading term of the actual expansion matches the upper bound.

From a coding perspective, this result implies that we should choose our Shannon random codebook to be sequences with mostly alternating 01 patterns and an occasional 00 pattern (i.e., occurs with probability $p_{00} \rightarrow 0$). It is also worth mentioning that this constraint costs nothing when the noise is large because the slope of the expansion matches the slope of the unconstrained BSC as $p_{00} \rightarrow 0$.

4.3 FSC Example: Intersymbol-Interference Channels in AWGN

Consider a family of finite-memory ISI channels parametrized by θ . Let the time- t output Y_t be a Gaussian whose mean is given by θ times a deterministic function of the current input and the previous k inputs. Under these conditions, the output process is a conditionally Gaussian HMP, with state $Q_t = (X_{t-1}, \dots, X_{t-k})$, as defined in Section 3.6. Moreover, the conditional entropy rate $H(\mathcal{Y}|\mathcal{X})$ only depends on the noise variance, which can be taken to be 1 without loss of generality. Therefore, θ -derivatives of the mutual information rate, $I(\mathcal{X}; \mathcal{Y}) = H(\mathcal{Y}) - H(\mathcal{Y}|\mathcal{X})$, depend only on θ -derivatives of the entropy rate $H(\mathcal{Y})$.

Let the mean of the output process induced by a state transition $Q_t = i$ to $Q_{t+1} = j$ be m_{ij} . One can explore the high-noise regime by keeping the noise variance fixed to 1 and letting $\theta \rightarrow 0$. In this case, one can combine (3.8) and Corollary 4.2 to see that

$$\mathcal{C}(\theta) = \frac{\theta^2}{2} \left[\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij}^2 - \left(\sum_{i,j \in \mathcal{Q}} \pi(i) p_{ij} m_{ij} \right)^2 \right] + o(\theta^2).$$

The first term in this expansion can be optimized over the input distribution p_{ij} , but there are a few caveats. Let $e_{ij} = \pi(i) p_{ij}$ be the edge occupancy probabilities that satisfy $\sum_{i,j \in \mathcal{Q}} e_{ij} = 1$, then

stationarity of the underlying Markov chain implies that $\sum_j (e_{ij} - e_{ji}) = 0$. One also finds that not all state transitions are valid, but setting $e_{ij} = 0$ if $(i, j) \notin \mathcal{V}$ gives the following convex⁴ optimization problem with linear constraints:

$$\begin{aligned} & \text{maximize} \sum_{i,j \in \mathcal{Q}} e_{ij} m_{ij}^2 - \left(\sum_{i,j \in \mathcal{Q}} e_{ij} m_{ij} \right)^2 \\ & \text{subject to} \sum_{i,j \in \mathcal{Q}} e_{ij} = 1 \\ & \sum_{j \in \mathcal{Q}} (e_{ij} - e_{ji}) = 0 \quad \forall i. \end{aligned}$$

A similar result is given in [41] for linear ISI channels with balanced inputs (i.e., a zero-mean input). In this case, the $\sum e_{ij} m_{ij}$ term is zero and the optimization problem is reduced to finding the maximum mean-weight cycle in a directed graph with edge weights m_{ij}^2 . The formula above generalizes the previous result to non-linear ISI channels and eliminates the zero-mean input requirement.

5 Connection to the Formula of Vontobel et al.

The results of this paper are closely related to an observation by Vontobel et al. [42] that the first part of generalized Blahut-Arimoto algorithm for FSCs actually computes the derivative of the mutual information. Their result is somewhat different because it considers derivatives with respect to the edge occupancy probabilities $\pi(i)p_{ij}$ rather than the observation probabilities. Their approach is also dissimilar because the answer is given exactly for finite blocks rather than focusing on the asymptotically long blocks and the forward/backward stationary measures. Moreover, the result in this paper does not apply to changes in the HMP which change the stationary distribution π of the while the derivative result in [42] focuses exclusively on changes in the edge occupancy probabilities.

Ideally, one would have a unified treatment of the derivative, with respect to changes in both the edge occupancy probabilities $\pi(i)p_{ij}$ and the observation probabilities, of the entropy rate of a FSC. Indeed, a simple formula, in terms of forward/backward stationary measures, can be cobbled together by translating the derivative formula in [42] to stationary measures and combining this with Theorem 3.2. To clarify the connection, their result is shown first in terms of conditional density functions for α and β . Paraphrasing their result, in terms of the derivative of the edge occupancy probabilities $\Delta_{ij} = \frac{d}{d\theta} \pi(i)p_{ij}|_{\theta=0}$, gives

$$\frac{d}{d\theta} H(\mathcal{X}|\mathcal{Y}; \theta)|_{\theta=0} = - \sum_{i,j \in \mathcal{Q}} \Delta_{ij} \int_{\mathcal{A}_0 \times \mathcal{A}_0} f_{\alpha|Q_t}(\alpha|i) f_{\beta|Q_{t+1}}(\beta|j) \sum_{y \in \mathcal{Y}} h_{ij}(y) \ln \frac{\alpha(i) M_{ij}(y) \beta(j)}{\sum_{k \in \mathcal{Q}} \alpha(i) M_{ik}(y) \beta(k)} d\alpha d\beta.$$

One can decompose this formula to see that the term Δ_{ij} gives the change in the edge occupancy probability, the term $f_{\alpha|Q_t}(\alpha|i) f_{\beta|Q_{t+1}}(\beta|j) f_{Y|Q_t Q_{t+1}}(y|i, j)$ is the probability of α, β, y given the transition, and the logarithmic term gives the contribution to $H(Q_{t+1} = j | Q_t = i, Y_{-\infty}^{\infty})$ for this α, β, y .

Next, we modify this expression to use unconditional α, β distributions with

$$\begin{aligned} \frac{d}{d\theta} H(\mathcal{X}|\mathcal{Y}; \theta)|_{\theta=0} & \stackrel{(a)}{=} - \sum_{i,j \in \mathcal{Q}} \Delta_{ij} \int_{\mathcal{A}_0 \times \mathcal{A}_0} \frac{\mu_i(d\alpha)}{\pi(i)} \cdot \frac{\nu_j(d\beta)}{\pi(j)} \sum_{y \in \mathcal{Y}} h_{ij}(y) \ln \frac{\alpha(i) M_{ij}(y) \beta(j)}{\sum_{k \in \mathcal{Q}} \alpha(i) M_{ik}(y) \beta(k)} \\ & \stackrel{(b)}{=} - \sum_{i,j \in \mathcal{Q}} \Delta_{ij} \int_{\mathcal{A}_0 \times \mathcal{A}_0} \frac{\mu(d\alpha) \alpha(i)}{\pi(i)} \cdot \frac{\nu(d\beta) \beta(j) \pi(j)}{\pi(j)} \sum_{y \in \mathcal{Y}} h_{ij}(y) \ln \frac{\alpha(i) M_{ij}(y) \beta(j)}{\sum_{k \in \mathcal{Q}} \alpha(i) M_{ik}(y) \beta(k)} \\ & \stackrel{(c)}{=} - \sum_{i,j \in \mathcal{Q}} \Delta_{ij} \int_{\mathcal{A}_0 \times \mathcal{A}_0} \mu(d\alpha) \nu(d\beta) \sum_{y \in \mathcal{Y}} \frac{\alpha(i) M_{ij}(y) \beta(j)}{\pi(i) p_{ij}} \ln \frac{\alpha(i) M_{ij}(y) \beta(j)}{\sum_{k \in \mathcal{Q}} \alpha(i) M_{ik}(y) \beta(k)}, \end{aligned}$$

⁴The objective function is actually concave, but one can negate the objective and minimize instead.

where (a) holds because $\frac{\mu_i(d\alpha)}{\pi(i)}$ is the conditional density of α given the true state is i and $\frac{\nu_j(d\beta)}{\pi(j)}$ is the conditional density of β given the true state is j , (b) follows from Lemmas 2.21 and 2.25, and (c) follows from $M_{ij}(y) = p_{ij}h_{ij}(y)$. Finally, using $H(\mathcal{Y};\theta) = H(\mathcal{X};\theta) - H(\mathcal{X}|\mathcal{Y};\theta) + H(\mathcal{Y}|\mathcal{X};\theta)$ and

$$\begin{aligned}\frac{d}{d\theta}H(\mathcal{X};\theta)|_{\theta=0} &= - \sum_{i,j \in \mathcal{Q}} \Delta_{ij} \ln p_{ij} \\ \frac{d}{d\theta}H(\mathcal{Y}|\mathcal{X};\theta)|_{\theta=0} &= - \sum_{i,j \in \mathcal{Q}} \Delta_{ij} \sum_{y \in \mathcal{Y}} h_{ij}(y) \ln h_{ij}(y),\end{aligned}$$

we find that $\frac{d}{d\theta}H(\mathcal{Y};\theta)|_{\theta=0}$ is given by

$$- \sum_{i,j \in \mathcal{Q}} \Delta_{ij} \int_{\mathcal{A}_0 \times \mathcal{A}_0} \mu(d\alpha) \nu(d\beta) \sum_{y \in \mathcal{Y}} \left[h_{ij}(y) \ln M_{ij}(y) + \frac{\alpha(i)M_{ij}(y)\beta(j)}{\pi(i)p_{ij}} \ln \frac{\alpha(i)M_{ij}(y)\beta(j)}{\sum_{k \in \mathcal{Q}} \alpha(i)M_{ik}(y)\beta(k)} \right].$$

It is straightforward to combine this Theorem 3.2, though the final expression is even more unwieldy.

6 Conclusions

This paper considers the derivative of the entropy rate for general hidden Markov processes and derives a closed-form expression for this derivative in high-noise limit. An application is presented relating to the achievable information rates of finite-state channels. Again, a closed-form expression is derived for the high-noise limit. Two examples of interest are considered. First, transmission over a BSC under a (0,1) RLL constraint is treated and the capacity-achieving input distribution is derived in the high-noise limit. Second, an intersymbol interference channel in AWGN is considered and the capacity is derived in the high-noise limit.

Acknowledgement. The author would like to thank an anonymous reviewer for catching a number of errors and inconsistencies in the paper. He is also grateful to Pascal Vontobel for his excellent comments on an earlier draft. This work also benefited from interesting discussions with Brian Marcus and is a natural extension of past work with Paul H. Siegel and Joseph B. Soriaga.

A Technical Details

A.1 Lemmas for Theorem 3.2

Lemma A.1. Consider function $F(\alpha, \beta) = -\alpha^T M' \beta \log(\alpha^T M \beta)$ where M is a non-negative matrix and M' is a real matrix. This function is Lipschitz continuous w.r.t. $\|\cdot\|_1$ on $(\alpha, \beta) \in \mathcal{P}_\delta \times \mathcal{B}_\delta$ where $\mathcal{B}_\delta = \{u \in \mathcal{A}_\delta \mid \pi^T u = 1\}$, $\eta = \min_i \pi(i) > 0$, and $\delta > 0$. This implies that

$$\begin{aligned}|F(\alpha, \beta) - F(\alpha', \beta)| &\leq L_\alpha \|\alpha - \alpha'\|_1 \\ |F(\alpha, \beta) - F(\alpha, \beta')| &\leq L_\beta \|\beta - \beta'\|_1 \\ |F(\alpha, \beta) - F(\alpha', \beta')| &\leq L_\alpha \|\alpha - \alpha'\|_1 + L_\beta \|\beta - \beta'\|_1,\end{aligned}$$

where $c = \delta^2 \sum_{i,j} M_{ij}$ and

$$\begin{aligned}L_\alpha &= \|M\|_1 \frac{1}{\eta} \log \frac{1}{c} + \|M'\|_1 \|M\|_1 \frac{1}{\eta^2 c} \\ L_\beta &= \|M\|_\infty \log \frac{1}{c} + \|M'\|_\infty \|M\|_\infty \frac{1}{c}.\end{aligned}$$

Proof. Let $G : \mathbb{R}^m \rightarrow \mathbb{R}$ be any function that is differentiable on a convex set $D \subseteq \mathbb{R}^m$. Then, the mean value theorem of vector calculus implies that

$$G(y) - G(x) = G'(x + t(y-x))^T (y-x)$$

for some $t \in [0, 1]$. Applying Hölder's inequality allows one to upper bound the Lipschitz constant w.r.t. $\|\cdot\|_1$ and gives the upper bound

$$\begin{aligned} G(y) - G(x) &\leq \sup_{t \in [0, 1]} \|G'(x + t(y - x))\|_\infty \|x - y\|_1 \\ &\leq \sup_{z \in D} \|G'(z)\|_\infty \|x - y\|_1. \end{aligned}$$

Since $F(\alpha, \beta)$ is differentiable w.r.t. α , we can bound the Lipschitz constant L_α with

$$\begin{aligned} L_\alpha &= \sup_{\alpha \in \mathcal{P}_\delta} \sup_{\beta \in \mathcal{B}} \sup_{\|u\|_\infty \leq 1} \left| u^T M' \beta \log \frac{1}{\alpha^T M \beta} - \alpha^T M' \beta \frac{u^T M \beta}{\alpha^T M \beta} \right| \\ &\stackrel{(a)}{\leq} \sup_{\alpha \in \mathcal{P}_\delta} \sup_{\beta \in \mathcal{B}} \sup_{\|u\|_\infty \leq 1} \left[|u^T M' \beta| \log \frac{1}{c} + |\alpha^T M' \beta| |u^T M \beta| \frac{1}{c} \right] \\ &\stackrel{(b)}{\leq} \|M\|_1 \|\beta\|_1 \log \frac{1}{c} + \|M'\|_1 \|\beta\|_1 \|M\|_1 \|\beta\|_1 \frac{1}{c} \\ &\stackrel{(c)}{\leq} \|M\|_1 \frac{1}{\eta} \log \frac{1}{c} + \|M'\|_1 \|M\|_1 \frac{1}{\eta^2 c}, \end{aligned} \tag{A.1}$$

where (a) follows from $\alpha^T M \beta \geq c$ with $c = \delta^2 \sum_{i,j} M_{ij}$, (b) follows from $|x^T M y| \leq \|x\|_\infty \|M\|_1 \|y\|_1$, and (c) follows from $\|\beta\|_1 \leq \eta^{-1}$ which holds because $\pi^T \beta = 1$.

Likewise $F(\alpha, \beta)$ is differentiable w.r.t. β and we can bound the Lipschitz constant L_β with

$$\begin{aligned} L_\beta &= \sup_{\alpha \in \mathcal{P}_\delta} \sup_{\beta \in \mathcal{B}} \sup_{\|u\|_\infty \leq 1} \left| \alpha^T M' u \log \frac{1}{\alpha^T M \beta} - \alpha^T M' \beta \frac{\alpha^T M u}{\alpha^T M \beta} \right| \\ &\stackrel{(a)}{\leq} \sup_{\alpha \in \mathcal{P}_\delta} \sup_{\beta \in \mathcal{B}} \sup_{\|u\|_\infty \leq 1} \left[|\alpha^T M' u| \log \frac{1}{c} + |\alpha^T M' \beta| |\alpha^T M u| \frac{1}{c} \right] \\ &\stackrel{(b)}{\leq} \|M\|_\infty \log \frac{1}{c} + \|M'\|_\infty \|M\|_\infty \frac{1}{c}, \end{aligned} \tag{A.2}$$

where (a) is the same as above and (b) follows from $|x^T M y| \leq \|x\|_1 \|M\|_\infty \|y\|_\infty$. \square

Lemma A.2. *If the HMP is ϵ -primitive for $\epsilon > 0$, then for some $\gamma < 1$ and $C < \infty$ we have*

$$\sum_{y \in \mathcal{Y}} E \left[\alpha^T M'(y) \beta \log (\alpha^T M(y) \beta) - \alpha_j^T M'(y) \beta_{j+1} \log (\alpha_j^T M(y) \beta_{j+1}) \right] \leq 2\bar{L}_\alpha C \gamma^{j-1} + 2\bar{L}_\beta C \gamma^{n-j+1},$$

where $c(y) = \delta^2 \sum_{i,j} [M(y)]_{ij}$ and

$$\begin{aligned} \bar{L}_\alpha &= \sum_{y \in \mathcal{Y}} \left[\|M(y)\|_1 \frac{1}{\eta} \log \frac{1}{c(y)} + \|M'(y)\|_1 \|M(y)\|_1 \frac{1}{\eta^2 c(y)} \right] \\ \bar{L}_\beta &= \sum_{y \in \mathcal{Y}} \left[\|M(y)\|_\infty \log \frac{1}{c(y)} + \|M'(y)\|_\infty \|M(y)\|_\infty \frac{1}{c(y)} \right]. \end{aligned}$$

The expectation assumes that α, β are drawn from their respective stationary distributions while α_j, β_{j+1} are drawn from the distributions implied by an arbitrary initialization of α_1, β_{n+1} .

Proof. Since the HMP is ϵ -primitive for $\epsilon > 0$, there is a δ such that $\min_i \alpha_i > \delta$ and $\min_i \beta_i > \delta$ on the entire support of α, β . It also follows that $\eta = \min_i \pi(i) > 0$. Now, consider the function $F_y(\alpha, \beta) = -\alpha^T M'(y) \beta \log (\alpha^T M(y) \beta)$. Under these conditions, Lemma A.1 shows that this function is Lipschitz continuous w.r.t. $\|\cdot\|_1$ on the support of α, β with Lipschitz constants $L_\alpha(y)$ and $L_\beta(y)$ defined

by generalizing (A.1) and (A.2). Therefore, we can write

$$\begin{aligned}
\sum_{y \in \mathcal{Y}} E_{\alpha, \beta} [F_y(\alpha, \beta) - F_y(\alpha_j, \beta_{j+1})] &\leq \sum_{y \in \mathcal{Y}} E_{\alpha, \beta} [L_\alpha(y) \|\alpha - \alpha_j\|_1 + L_\beta(y) \|\beta - \beta_{j+1}\|_1] \\
&\stackrel{(a)}{\leq} \sum_{y \in \mathcal{Y}} E_{\alpha, \beta} [L_\alpha(y) 2d(\alpha, \alpha_j) + L_\beta(y) 2d(\beta, \beta_{j+1})] \\
&\stackrel{(b)}{\leq} \sum_{y \in \mathcal{Y}} L_\alpha(y) 2C\gamma^{j-1} + \sum_{y \in \mathcal{Y}} L_\beta(y) 2C\gamma^{n-j+1},
\end{aligned}$$

where (a) follows from Lemma 2.4 and (b) follows from Lemma 2.12 because the HMP is ϵ -primitive. \square

A.2 Proof of Lemma 3.4

Proof. The first two results follow from Lemmas 2.21 and 2.25. Substituting and integrating gives and

$$\int_{\mathcal{A}_0} \mu(d\alpha) \alpha(q) = \int_{\mathcal{A}_0} \underbrace{\mu_q(d\alpha)}_{\Pr(Q=q, \alpha \in d\alpha)} = \Pr(Q = q)$$

and

$$\int_{\mathcal{A}_0} \nu(d\beta) \beta(q) = \int_{\mathcal{A}_0} \frac{1}{\pi(q)} \underbrace{\nu_q(d\beta)}_{\Pr(Q=q, \beta \in d\beta)} = 1.$$

Using the fact that

$$\sum_{y \in \mathcal{Y}} M(y) = P,$$

we can evaluate the third and fourth results with

$$\int_{\mathcal{A}_0} \mu(d\alpha) \alpha^T \sum_{y \in \mathcal{Y}} M(y) \beta = \pi^T P \beta = \pi^T \beta = 1$$

and

$$\alpha^T \sum_{y \in \mathcal{Y}} M(y) \int_{\mathcal{A}_0} \nu(d\beta) \beta = \alpha^T P \mathbf{1} = \alpha^T \mathbf{1} = 1.$$

Finally, the fifth result follows from

$$\frac{d}{d\theta} \int_{\mathcal{A}_0} \mu(d\alpha) \alpha^T \sum_{y \in \mathcal{Y}} M_\theta(y) \int_{\mathcal{A}_0} \nu(d\beta) \beta = \frac{d}{d\theta} \pi^T P_\theta \mathbf{1} = \frac{d}{d\theta} 1 = 0.$$

\square

A.3 Proof of Theorem 3.6

Proof. First, we point out that $\lim_{\theta \rightarrow \theta^*} M_\theta(y) = s(y) P$ implies that output symbols provide no state information at $\theta = \theta^*$ so that $H(\mathcal{Y}; \theta^*) = H(Y_1; \theta^*)$. This also implies that, at $\theta = \theta^*$, the forward and backward Blackwell measures are Dirac measures, $\mu(A) = \mathbb{1}_A(\pi)$ and $\nu(B) = \mathbb{1}_B(\mathbf{1})$, concentrated on $\pi, \mathbf{1}$. By Theorem 3.2, the derivative of the entropy rate is uniformly continuous on D and we have

$$\begin{aligned}
\lim_{\theta \rightarrow \theta^*} \frac{d}{d\theta} H(\mathcal{Y}; \theta) &= - \lim_{\theta \rightarrow \theta^*} E_{\alpha, \beta} \left[\sum_{y \in \mathcal{Y}} \alpha^T M'_\theta(y) \beta \ln(\alpha^T M_\theta(y) \beta) \right] \\
&= - \sum_{y \in \mathcal{Y}} \pi^T M'(y) \mathbf{1} \ln(s(y)) - \pi^T \left(\sum_{y \in \mathcal{Y}} M'(y) \right) \mathbf{1} \ln(\pi^T P \mathbf{1}) \\
&\stackrel{(a)}{=} - \sum_{y \in \mathcal{Y}} \pi^T M'(y) \mathbf{1} \ln(s(y)),
\end{aligned}$$

where (a) holds because $\pi^T P \mathbf{1} = 1$.

For the 2nd derivative, we apply the derivative shortcut a second time by noting that

$$g_n''(\theta) = \sum_{i=1}^n \sum_{j=1}^n \frac{\partial}{\partial \theta_i} \frac{\partial}{\partial \theta_j} g_n(\theta).$$

Applying this to $g_n(\theta_1, \dots, \theta_n)$ for the entropy rate gives $g_n''(\theta^*)$

$$\begin{aligned} &= -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \frac{\partial}{\partial \theta_i} \frac{\partial}{\partial \theta_j} \sum_{y_1^n \in \mathcal{Y}^n} \pi^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \cdot \log \left[\pi^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \right] \Big|_{(\theta_1, \dots, \theta_n) = (\theta^*, \dots, \theta^*)} \\ &\stackrel{(a)}{=} -\frac{1}{n} \sum_{i=1}^n \frac{\partial}{\partial \theta_i} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \pi^T M(y_1^{j-1}) M'_{\theta_j}(y_j) M(y_{j+1}^n) \mathbf{1} \cdot \log \left[\pi^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \right] \Big|_{(\theta_1, \dots, \theta_n) = (\theta^*, \dots, \theta^*)} \\ &\quad - \frac{1}{n} \sum_{i=1}^n \frac{\partial}{\partial \theta_i} \sum_{j=1}^n \frac{\partial}{\partial \theta_j} \sum_{y_1^n \in \mathcal{Y}^n} \pi^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \Big|_{(\theta_1, \dots, \theta_n) = (\theta^*, \dots, \theta^*)} \tag{A} \\ &= -\frac{1}{n} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \pi^T M(y_1^{j-1}) M''(y_j) M(y_{j+1}^n) \mathbf{1} \cdot \log \left[\pi^T \left(\prod_{t=1}^n M(y_t) \right) \mathbf{1} \right] \tag{T1} \\ &\quad - \frac{1}{n} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \frac{\left(\pi^T M(y_1^{j-1}) M'(y_j) M(y_{j+1}^n) \mathbf{1} \right)^2}{\pi^T \left(\prod_{t=1}^n M(y_t) \right) \mathbf{1}} \tag{T2} \\ &\quad - \frac{2}{n} \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_1^n \in \mathcal{Y}^n} \pi^T M(y_1^{i-1}) M'(y_i) M(y_{i+1}^{j-1}) M'(y_j) M(y_{j+1}^n) \mathbf{1} \cdot \log \left[\pi^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \right] \tag{T3} \\ &\quad - \frac{2}{n} \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_1^n \in \mathcal{Y}^n} \frac{\left(\pi^T M(y_1^{i-1}) M'(y_i) M(y_{i+1}^n) \mathbf{1} \right) \left(\pi^T M(y_1^{j-1}) M'(y_j) M(y_{j+1}^n) \mathbf{1} \right)}{\pi^T \left(\prod_{t=1}^n M(y_t) \right) \mathbf{1}}, \tag{T4} \end{aligned}$$

where the term labeled (A) is zero because it equals $-\frac{1}{n} \frac{d^2}{d^2 \theta} 1$. Using the term labels in the equation (i.e., T1, T2, ...), we see that $g_n''(\theta^*) = T_1 + T_2 + T_3 + T_4$, where the terms T_1, T_2 are associated with $i = j$, and the terms T_3, T_4 are associated with $i \neq j$. Using this decomposition, we can reduce each term separately.

For the first term, $M(y) = s(y)P$ implies that

$$\begin{aligned} T_1 &= -\frac{1}{n} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \pi^T M(y_1^{j-1}) M''(y_j) M(y_{j+1}^n) \mathbf{1} \cdot \log \left[\pi^T \left(\prod_{t=1}^n M(y_t) \right) \mathbf{1} \right] \\ &= -\frac{1}{n} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \frac{s(y_1^n)}{s(y_j)} \pi^T M''(y_j) \mathbf{1} \cdot \log(s(y_1^n)) \\ &= -\frac{1}{n} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \left(\frac{s(y_1^n)}{s(y_j)} \pi^T M''(y_j) \mathbf{1} \cdot \log(s(y_j)) + \frac{s(y_1^n)}{s(y_j)} \pi^T M''(y_j) \mathbf{1} \sum_{k=1, k \neq j}^n \log(s(y_k)) \right) \\ &\stackrel{(a)}{=} -\frac{1}{n} \sum_{j=1}^n \left(\sum_{y_j \in \mathcal{Y}} \pi^T M''(y_j) \mathbf{1} \cdot \log(s(y_j)) + 0 \right) \\ &= -\sum_{y \in \mathcal{Y}} \pi^T M''(y) \mathbf{1} \cdot \log(s(y)), \end{aligned}$$

where (a) follows from the fact that

$$\sum_{y_j \in \mathcal{Y}} \frac{s(y_1^n)}{s(y_j)} \pi^T M''(y_j) \mathbf{1} \sum_{k=1, k \neq j}^n \log(s(y_k)) = \left(\prod_{i=1, i \neq j}^n s(y_i) \right) \left(\sum_{k=1, k \neq j}^n \log(s(y_k)) \right) \sum_{y_j \in \mathcal{Y}} \pi^T M''(y_j) \mathbf{1} = 0.$$

For the second term, $M(y) = s(y)P$ implies that

$$\begin{aligned} T_2 &= -\frac{1}{n} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \frac{\left(\pi^T M(y_1^{j-1}) M'(y_j) M(y_{j+1}^n) \mathbf{1} \right)^2}{\pi^T \left(\prod_{t=1}^n M(y_t) \right) \mathbf{1}} \\ &= -\frac{1}{n} \sum_{j=1}^n \sum_{y_1^n \in \mathcal{Y}^n} \frac{s(y_1^n)^2}{s(y_1^n) s(y_j)^2} (\pi^T M'(y_j) \mathbf{1})^2 \\ &= -\frac{1}{n} \sum_{j=1}^n \sum_{y_j \in \mathcal{Y}} \frac{(\pi^T M'(y_j) \mathbf{1})^2}{s(y_j)} \\ &= -\sum_{y \in \mathcal{Y}} \frac{(\pi^T M'(y) \mathbf{1})^2}{s(y)} \\ &= -\sum_{y \in \mathcal{Y}} \frac{(\pi^T M'(y) \mathbf{1})^2}{\pi^T M(y) \mathbf{1}} \end{aligned}$$

For the third term, we notice first that $\sum_{y \in \mathcal{Y}} M'(y) = 0$ implies

$$\sum_{y_i, y_j, y_k \in \mathcal{Y}^n} \pi^T M'(y_i) P^{j-i-1} M'(y_j) \mathbf{1} \cdot \log(s(y_k)) = 0$$

if either $i \neq k$ or $j \neq k$. This gives

$$\begin{aligned} T_3 &= -\frac{2}{n} \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_1^n \in \mathcal{Y}^n} \pi^T M(y_1^{i-1}) M'(y_i) M(y_{i+1}^{j-1}) M'(y_j) M(y_{j+1}^n) \mathbf{1} \cdot \log \left[\pi^T \left(\prod_{t=1}^n M_{\theta_t}(y_t) \right) \mathbf{1} \right] \\ &= -\frac{2}{n} \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_1^n \in \mathcal{Y}^n} \frac{s(y_1^n)}{s(y_i) s(y_j)} \pi^T M'(y_i) P^{j-i-1} M'(y_j) \mathbf{1} \cdot \log(s(y_1^n)) \\ &= -\frac{2}{n} \sum_{k=1}^n \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_i, y_j, y_k \in \mathcal{Y}^n} \pi^T M'(y_i) P^{j-i-1} M'(y_j) \mathbf{1} \cdot \log(s(y_k)) \\ &= 0 \end{aligned}$$

because $i < j$.

For the fourth term, we have

$$\begin{aligned} T_4 &= -\frac{2}{n} \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_1^n \in \mathcal{Y}^n} \frac{(\pi^T M(y_1^{i-1}) M'(y_i) M(y_{i+1}^n) \mathbf{1}) (\pi^T M(y_1^{j-1}) M'(y_j) M(y_{j+1}^n) \mathbf{1})}{\pi^T \left(\prod_{t=1}^n M(y_t) \right) \mathbf{1}} \\ &= -\frac{2}{n} \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_1^n \in \mathcal{Y}^n} \frac{s(y_1^n)^2}{s(y_1^n) s(y_i) s(y_j)} (\pi^T M'(y_i) \mathbf{1}) (\pi^T M'(y_j) \mathbf{1}) \\ &= -\frac{2}{n} \sum_{j=1}^n \sum_{i=1}^{j-1} \sum_{y_i, y_j \in \mathcal{Y}^n} (\pi^T M'(y_i) \mathbf{1}) (\pi^T M'(y_j) \mathbf{1}) \\ &= 0 \end{aligned}$$

because $\sum_{y \in \mathcal{Y}} M'(y) = 0$. □

References

- [1] D. Arnold and H. Loeliger. On the information rate of binary-input channels with memory. In *Proc. IEEE Int. Conf. Commun.*, pages 2692–2695, Helsinki, Finland, June 2001.
- [2] D. Arnold, H. A. Loeliger, P. O. Vontobel, A. Kavčić, and W. Zeng. Simulation-based computation of information rates for channels with memory. *IEEE Trans. Inform. Theory*, 52(8):3498–3508, Aug. 2006.
- [3] D. M. Arnold. *Computing information rates of finite-state models with application to magnetic recording*. PhD thesis, Swiss Federal Institute of Technology, Zurich, 2003.
- [4] R. G. Bartle and D. R. Sherbert. *Introduction to Real Analysis*. Wiley, 3rd edition, 1999.
- [5] L. E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *Ann. Math. Stats.*, 37:1554–1563, Dec. 1966.
- [6] P. Billingsley. *Convergence of probability measures*. Wiley, 2nd edition, 1999.
- [7] J. Birch. Approximations for the entropy for functions of Markov chains. *Ann. Math. Stats.*, 33(3):930–938, Sept. 1962.
- [8] D. Blackwell. Entropy of functions of finite-state Markov chains. *Trans. First Prague Conf. on Inform. Theory, Stat. Dec. Fun., Rand. Processes*, pages 13–20, 1957.
- [9] D. Blackwell, L. Breiman, and A. J. Thomasian. Proof of Shannon’s transmission theorem for finite-state indecomposable channels. *Ann. Math. Stats.*, 29:1209–1220, Dec. 1958.
- [10] J. Chen and P. H. Siegel. Markov processes asymptotically achieve the capacity of finite-state intersymbol interference channels. *IEEE Trans. Inform. Theory*, 54(3):1295–1303, 2008.
- [11] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, 1991.
- [12] E. Domany, I. Kanter, O. Zuk, and M. Aizenman. From finite-system entropy to entropy rate for a hidden Markov process. *IEEE Signal Processing Letters*, 13(9):517–520, Sept. 2006.
- [13] Y. Ephraim and N. Merhav. Hidden Markov processes. *IEEE Trans. Inform. Theory*, 48(6):1518–1569, June 2002.
- [14] H. Furstenberg and H. Kesten. Products of random matrices. *Ann. Math. Stats.*, 31:457–469, June 1960.
- [15] H. Furstenberg and Y. Kifer. Random matrix products and measures on projective spaces. *Israel Journal of Mathematics*, 46(1):12–32, 1983.
- [16] R. G. Gallager. *Information Theory and Reliable Communication*. Wiley, New York, NY, USA, 1968.
- [17] A. J. Goldsmith and P. P. Varaiya. Capacity, mutual information, and coding for finite-state Markov channels. *IEEE Trans. Inform. Theory*, 42(3):868–886, May 1996.
- [18] G. Han and B. Marcus. Analyticity of entropy rate of hidden Markov chains. *IEEE Trans. Inform. Theory*, 52(12):5251–5266, Dec. 2006.
- [19] G. Han and B. Marcus. Asymptotics of noisy constrained channel capacity. In *Proc. IEEE Int. Symp. Information Theory*, pages 991–995, Nice, France, June 2007.
- [20] G. Han and B. Marcus. Derivatives of entropy rate in special families of hidden Markov chains. *IEEE Trans. Inform. Theory*, 53(7):2642–2652, 2007.

- [21] T. E. Harris. On chains of infinite order. *Pacific J. Math*, 5(1):707–724, 1955.
- [22] T. Holliday, A. Goldsmith, and P. Glynn. Capacity of finite state channels based on Lyapunov exponents of random matrices. *IEEE Trans. Inform. Theory*, 52(8):3509–3532, Aug. 2006.
- [23] F. Jelinek. Continuous speech recognition by statistical methods. *Proc. of the IEEE*, 64(4):532–556, 1976.
- [24] A. Kavčić. On the capacity of Markov sources over noisy channels. In *Proc. IEEE Global Telecom. Conf.*, pages 2997–3001, San Antonio, Texas, USA, Nov. 2001.
- [25] A. Krogh, M. Brown, I. Mian, K. Sjolander, and D. Haussler. Hidden Markov models in computational biology: Applications to protein modeling. *J. Molecular Biology*, 235(5):1501–1531, 1994.
- [26] F. Le Gland and L. Mevel. Basic properties of the projective product with application to products of column-allowable nonnegative matrices. *Math. Control Signals Systems*, 13(1):41–62, July 2000.
- [27] F. Le Gland and L. Mevel. Exponential forgetting and geometric ergodicity in hidden Markov models. *Math. Control Signals Systems*, 13(1):63–93, July 2000.
- [28] C. Méasson, A. Montanari, T. J. Richardson, and R. L. Urbanke. Life above threshold: From list decoding to area theorem and MSE. *Arxiv preprint cs.IT/0410028*, 2004.
- [29] C. Méasson, A. Montanari, and R. L. Urbanke. Maxwell construction: The hidden bridge between iterative and maximum a posteriori decoding. *IEEE Trans. Inform. Theory*, 54(12):5277–5307, Dec. 2008.
- [30] M. Mushkin and I. Bar-David. Capacity and coding for Gilbert-Elliott channels. *IEEE Trans. Inform. Theory*, 35(6):1277–1290, Nov. 1989.
- [31] E. Ordentlich and T. Weissman. New bounds on the entropy rate of hidden Markov processes. In *Proc. IEEE Inform. Theory Workshop*, pages 117–122, San Antonio, TX, Oct. 2004.
- [32] E. Ordentlich and T. Weissman. Approximations for the entropy rate of a hidden Markov process. In *Proc. IEEE Int. Symp. Information Theory*, pages 2198–2202, Adelaide, Australia, Sept. 2005.
- [33] E. Ordentlich and T. Weissman. On the optimality of symbol-by-symbol filtering and denoising. *IEEE Trans. Inform. Theory*, 52(1):19–40, Jan. 2006.
- [34] V. I. Oseledec. A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems. *Trans. Moscow Math. Soc.*, pages 197–231, 1968.
- [35] Y. Peres. Domains of analytic continuation for the top Lyapunov exponent. *Ann. Inst. H. Poincaré Probab. Statist.*, 28(1):131–148, 1992.
- [36] T. Petrie. Probabilistic functions of finite state Markov chains. *Ann. Math. Stats.*, 40(1):97–115, Feb. 1969.
- [37] H. D. Pfister. *On the Capacity of Finite State Channels and the Analysis of Convolutional Accumulate-m Codes*. PhD thesis, University of California, San Diego, La Jolla, CA, USA, March 2003.
- [38] H. D. Pfister, J. B. Soriaga, and P. H. Siegel. On the achievable information rates of finite state ISI channels. In *Proc. IEEE Global Telecom. Conf.*, pages 2992–2996, San Antonio, Texas, USA, Nov. 2001.
- [39] D. Ruelle. Analyticity properties of the characteristic exponents of random matrix products. *Adv. Math.*, 32:68–80, 1979.
- [40] E. Seneta. *Non-Negative Matrices: An Introduction to Theory and Applications*. Wiley, New York, NY, USA, 2nd edition, 1981.

- [41] J. B. Soriaga, H. D. Pfister, and P. H. Siegel. On the low-rate Shannon limit for binary intersymbol interference channels. *IEEE Trans. Commun.*, 51(12):1962–1964, Dec. 2003.
- [42] P. O. Vontobel, A. Kavčić, D. M. Arnold, and H. A. Loeliger. A generalization of the Blahut–Arimoto algorithm to finite-state channels. *IEEE Trans. Inform. Theory*, 54(5):1887–1918, 2008.
- [43] A. Ziv. Relative distance – an error measure in round-off error analysis. *Math. Comp.*, 39(160):563–569, Oct. 1982.
- [44] O. Zuk, I. Kanter, and E. Domany. The entropy of a binary hidden Markov process. *J. Stat. Phys.*, 121(3):343–360, Nov. 2005.